

# Testing An Earthquake Early Warning Algorithm For New Zealand

By

Rasika Nandana Walakulu Arachchige

A Thesis

submitted to the Victoria University of Wellington  
in partial fulfilment of the requirements for the

degree of Doctor of Philosophy



Te Herenga Waka - Victoria University of Wellington

2024



## Dedication

To all my teachers,  
including my first teachers, അമ്മ (Mom) and അച്ഛൻ (Dad),

And

to the late Kelly Savage

***Makere ana ngā here.***



# Abstract

Earthquake early warning systems (EEW) are valuable tools for minimising the impact of seismic events on communities and infrastructure. While some countries, like Japan, Taiwan, South Korea, Mexico, and the United States, have successfully implemented EEW systems, New Zealand has yet to establish a national EEW system. This thesis explores/evaluates a potential EEW algorithm for a national EEW system for New Zealand.

PLUM (Propagation of Local Undamped Motion) is a ground motion based EEW algorithm that offers several advantages over point source and finite fault EEW algorithms due to its robustness, speed, and simplicity. It does not rely on an earthquake source model, making it less computationally expensive. The PLUM algorithm has been successfully implemented in Japan and tested on United States West Coast data.

To evaluate the PLUM algorithm's performance and usability in the New Zealand context, it is necessary to test it using a range of earthquakes that could happen over the next century. Because there is insufficient recorded data from past damaging earthquakes, we used a 220000-year-long physics-based synthetic earthquake catalogue of New Zealand and generated synthetic seismograms using a ground motion simulator. We also analysed the latency of the current New Zealand GeoNet sensor network since the latency reduces the warning time of alerts.

We focused our evaluation on the Canterbury and Wellington regions. The Canterbury region has experienced several damaging earthquakes in the past 15 years, and a dense network of seismic instruments monitors part of it. The Wellington region is affected by a very complex set of faults, is considered very earthquake-prone and has a population centre adjacent to Cook Strait. The algorithm was tested for the Canterbury and Wellington regions in terms of alert correctness and timeliness, algorithm configuration settings, and the impact of some sensor failures.

The results showed that the PLUM algorithm works well in highly populated areas of Canterbury, but not as effectively in some southern parts of Wellington, due to the lack of stations in Cook Strait to provide warnings for earthquakes centred south of Wellington.

However, it does not perform well in rural areas of either region because the sensor density is too low, even with an increased warning radius.

For the highly populated areas of Canterbury, PLUM could provide Correct Timely Alerts (warning time > 0 sec) for more than 90% of the expected shakings in a 100-year interval. Low-populated regions with a low density of sensors received fewer Correct Timely Alerts. The most populated areas of Wellington could also receive more than 70% Correct Timely Alerts. However, PLUM also generated a large number of precautionary alerts (an alert is received, but only weak shaking is felt), which would need to be addressed. The results suggested an appropriate choice of PLUM warning radius in the New Zealand context would be more than the original 30 km radius designed for the Japan context. The analysis also demonstrates the robustness of the algorithm under simulated sensor or communication breakdowns, particularly in regions with high sensor density.

The thesis shows that PLUM could be a valuable component of a New Zealand EEW system, and the results will be beneficial for designing future earthquake early warning systems for New Zealand. The testing framework in the thesis would also be useful for testing the suitability of earthquake early warning algorithms using synthetic earthquake catalogues.

# Acknowledgements

I am deeply grateful to everyone who has supported and guided me throughout the journey of completing this thesis. This has not just been a thesis but a journey full of challenges, experiences, highs and lows, and growth.

First and foremost, I would like to express my sincere gratitude to my supervisors, Martha Savage, Caroline Holden, and Peter Andreae (Pondy), for their unwavering support, insightful guidance, and valuable feedback throughout this research process. Meeting with Pondy whenever needed felt like visiting a trusted doctor. During these sessions, I eagerly absorbed all the advice and remedies offered to address my challenges. Thank you, Pondy, for always being willing to listen to my academic and personal concerns. I am incredibly grateful to Martha and Caroline for being encouraging mentors who believed in me from the very beginning. Your support has been a constant source of motivation and confidence for me.

I would also like to thank my thesis examiners—Professor John Townend (School of Geography, Environment and Earth Sciences, Victoria University of Wellington), Dr. Jennifer Andrews (Seismologist, GNS Science), and Associate Professor Masumi Yamada (Disaster Prevention Research Institute, Kyoto University, Japan)—for their constructive comments and thoughtful suggestions that significantly improved the quality of this thesis.

I am also grateful to John Townend and Calum Chamberlain for their thoughtful feedback on my initial proposal, which significantly helped refine the scope and direction of this research. Special thanks to Richard Arnold for his invaluable advice on statistical analysis. I also wish to thank Sarah-Jayne McCurrach for her early supervision and guidance during the formative stages of this research.

This research was made possible through funding from several generous sources. I am thankful for the financial support provided by the Open University of Sri Lanka (OUSL) through the AHEAD program, the Wellington Doctoral Scholarship of Victoria University of Wellington (VUW), and the New Zealand Resilience to Nature's Challenges program (RNC2) under the "Earthquakes and Tsunami" theme.

The input and resources especially synthetic catalogue data, provided by the RNC2 program have been invaluable. I am particularly grateful for the insights from other program members, including Andy Howell, Camila Penney, and Laura Hughes which have greatly enriched this research. I would also like to acknowledge Bill Fry and Andy Nicol, the project leaders, for their leadership.

I am grateful to Jen Andrews for her feedback on the research, and to Julia Becker for her input on social science aspects. My thanks also go to Daniel Hill and Kevin Fenaughty from the National Emergency Management Agency for their insights into agency requirements.

My initial discussions with GeoNet were invaluable in understanding its architecture. Thank you to Richard Guest, Elisabetta D'Anastasio, and the other GeoNet staff for their support..I acknowledge the New Zealand GeoNet programme and its sponsors NHC, GNS Science, LINZ, NEMA and MBIE for providing data used in this study. I also appreciate the comments and feedback received from various researchers during conferences and the authors of the research papers I referred to throughout my studies.

I would like to extend my gratitude to Raj Prasanna, David Johnston, and the JCDR team from the Massey University low-cost EEW project, where I worked as a research officer before transitioning to this project. Your mentorship laid a foundation for my growth in this field.

I am grateful to Patricia Stein and Catherine Shone (student advisors), Barry Lewis (scholarships and doctoral admissions), Erica Schouten (student advocate, VUWSA), Miniruwani Samarakoon from PGSA, Gillian Turner, Neil Dodgson, and the staff at the Faculty of Science/Engineering and the Faculty of Graduate Research (FGR) for their continuous support and guidance. A big thank you also goes to Diana Siwiak , Tony McLoughlin, Morgan Holschier, Aleksandr Beliaev, Emily Brook, Colin Wilson and the Heads of Departments and staff at ECS and SGEES for their assistance throughout my PhD journey. I would also like to thank John Barclay and Edmund Salem from Maori Ora, and Louise Wanoa (FGR offshore support coordinator) and the staff at VUW International/Student Experience and Wellbeing for their support.

I am also grateful to Upeka Premarathne, my Master's supervisor, and Ravi de Mel, my undergraduate supervisor, for their guidance, as well as to my previous supervisors, Winston and Alvin. I am grateful to my friends Pushpakumara, Yamuna, Bimsara, Dilshan, Nuwan and Madushika for their support. Special thanks to the staff at OUSL and AHEAD, including Ariadurai, Rathnaweera, Bandunee, Wijesundara, Ruminda, Iresha, Nelaka, Gayani, Madhubhashini, Yasas, Tharangani and Samarasinghe, for their assistance.

I am thankful to my VUW PhD friends, Chethana, Erandi, and Hiroshika, for their friendship and support. I would also like to thank my Sri Lankan friends in Wellington—Achala, Chamila, Sudam, Ruwan, Chaminda, Amila, Suranga, Asela, Laksiri, Vajira, Awanthi, Clement, Jero, Sarda, Buddhika, Chanthujan, Nimrekha, Kavi, Chanaka, Wasana, and many others—who have been like an extended family to me here.

I want to express my gratitude to my colleagues from the VUW SL Student Association including Hesadi and the QuakeCore Emerging Research Chapter for their collaboration. Thanks also to my officemates in Maru at ECS and SGEES and the members of my research group, including El, Stephen, Carmen, Hassan, Laura and Jesse.

Finally, I am profoundly grateful to my wife, Wasanthi Gamage, and my little daughter, Tenuki, for their endless patience, love, and support. We missed each other dearly during the more than two years apart. I am deeply grateful to my mother, Gunawathi Wickramasinghe; my

father, Gnanawimala Walakulu Arachchige; and my two brothers, Isuru and Sameera, for their unwavering encouragement throughout this journey.

I also want to acknowledge the tools and resources that were instrumental in this research: EXSIM, Python (including ObsPy, Pandas, NumPy, and GeoPandas Matplotlib), MATLAB, and GMT (The Generic Mapping Tools), as well as Grammarly. Thank you all for your support and encouragement.

# Table of Contents

Abstract.....	iii
Acknowledgements.....	v
Table of Contents.....	viii
List of Figures .....	xi
List of Tables .....	xiv
Chapter 1 Introduction.....	1
1.1 Research goals and objectives .....	3
1.2 Organisation of the thesis .....	4
Chapter 2 Background and Literature Review .....	5
2.1 New Zealand context for an EEW system .....	5
2.2 Earthquake early warning algorithms .....	9
2.2.1 Point source algorithms .....	9
2.2.2 Finite fault algorithms.....	12
2.2.3 Ground motion based algorithms.....	15
2.2.4 On-site algorithms .....	19
2.2.5 Selection of an algorithm to test for NZ .....	20
2.3 Earthquake early warning systems .....	22
2.4 GeoNet sensor network and capabilities .....	23
2.5 EXSIM ground motion simulator .....	27
2.6 RNC2 synthetic earthquake catalogue.....	30
Chapter 3 Preparatory Analysis .....	33
3.1 Data latency of GeoNet sensor network.....	33
3.2 Initial exploration of PLUM on two recent NZ earthquakes. ....	37
3.3 Evaluating EXSIM on NZ earthquake data.....	42
3.3.1 Synthetic seismograms of Darfield earthquake.....	42
3.3.2 Comparison of the synthetic and GeoNet data .....	45
3.4 Testing PLUM with 209 synthetic earthquakes in the Canterbury region.....	54
3.4.1 Performance for different warning radii. ....	57
3.4.2 Adding new stations.....	62
3.4.3 Population density and PLUM for Canterbury region .....	69
3.5 Conclusions and limitations .....	71
Chapter 4 Evaluating PLUM on the Canterbury Region.....	75
4.1 Introduction.....	75
4.1.1 Noise mitigation for PLUM in Japan and USA.....	76
4.2 Methodology .....	79
4.2.1 Region of interest.....	80
4.2.2 Site classification .....	82

4.2.3	Random selection of sub-catalogues from the main catalogue .....	83
4.2.4	Synthetic seismogram generation for the catalogue events .....	89
4.2.5	Noise and error mitigation .....	91
4.2.6	PLUM algorithm .....	95
4.2.7	EEW performance indicators .....	99
4.2.8	Population-weighted measures .....	103
4.2.9	Assessment of the EEW performance in simulated failure scenarios .....	105
4.3	Results for the Canterbury region .....	107
4.3.1	Results for the Canterbury region by location .....	108
4.3.2	Results for the Canterbury region by population .....	121
4.3.3	Results for failure scenarios .....	132
4.4	Discussion .....	148
Chapter 5 Evaluating PLUM on the Wellington Region .....		155
5.1	Introduction .....	155
5.2	Methodology .....	156
5.2.1	Region of interest .....	156
5.2.2	Site classification .....	159
5.2.3	Selection of sub-catalogues from the main catalogue. ....	164
5.2.4	Synthetic seismogram generation for the catalogue events .....	166
5.2.5	PLUM algorithm .....	166
5.2.6	Performance Indicators .....	170
5.3	Results .....	170
5.3.1	Results for the Wellington region by location. ....	170
5.3.2	Results for the Wellington region by population. ....	180
5.3.3	Analysis of alerts for south Wellington City from Cook Strait earthquakes. ....	187
5.4	Discussion .....	188
Chapter 6 Conclusions and Future Work .....		193
6.1	Discussion and recommendations .....	193
6.2	Future work .....	200
References .....		207



# List of Figures

Figure 2-1 The National Warning System for Earthquakes and Tsunami .....	7
Figure 2-2 Plum algorithm .....	16
Figure 2-3 Alpha algorithm .....	19
Figure 2-4 GeoNet seismic and GPS sensor map (North and South Islands) .....	25
Figure 2-5 Distribution of nearest-station distances among 344 GeoNet strong motion stations .....	26
Figure 2-6 Typical GeoNet seismic data flow .....	27
Figure 2-7 Sample catalogue earthquakes from a 500-year period. ....	31
Figure 3-1 Strong motion sensor network - median latency .....	35
Figure 3-2 Distribution of the median latency of the strong motion sensors .....	36
Figure 3-3 Basic PLUM configuration .....	38
Figure 3-4 Three stations for PLUM simulation with the 2016 Kaikōura earthquake .....	39
Figure 3-5 Three stations for PLUM simulation with the 2013 Lake Grassmere earthquake .....	41
Figure 3-6 Darfield earthquake fault segments and slip distribution – fault model .....	45
Figure 3-7 Comparison of Synthetic and GeoNet PGA data .....	46
Figure 3-8 Comparison of modeled (synthetic) and observed (GeoNet) PGA Values. ....	47
Figure 3-9 Comparison of Synthetic and GeoNet detection time data .....	48
Figure 3-10 Comparison of the Synthetic and GeoNet seismograms: DFHS station. ....	50
Figure 3-11 Comparison of the Synthetic and GeoNet seismograms: CACS station.....	51
Figure 3-12 Comparison of the Synthetic and GeoNet seismograms: DORC station.....	52
Figure 3-13 PGA of Synthetic seismograms at grid points across Canterbury.....	53
Figure 3-14 GeoNet Shaking Layers Map for the Darfield Earthquake.....	54
Figure 3-15 Canterbury region with grid.....	55
Figure 3-16 Fraction of Correct Timely Alerts for 209 synthetic earthquakes.....	59
Figure 3-17 Fraction of False Alerts for 209 synthetic earthquakes .....	60
Figure 3-18 Total of CTA, MLA and FA at different PLUM radii. ....	61
Figure 3-19 Five new station locations.....	64
Figure 3-20 $CTA_{norm}$ for warning radius 30 km after adding 5 stations.....	65
Figure 3-21 Optimal position for a single new station .....	68
Figure 3-22 Population density of Canterbury region .....	70
Figure 4-1 JMA intensity filter .....	77
Figure 4-2 The selected Canterbury region with $0.1^\circ \times 0.1^\circ$ grid.....	81
Figure 4-3 Canterbury region with sensors .....	82
Figure 4-4 Synthetic earthquakes .....	84
Figure 4-5 Historical catalogue of shallow earthquakes (M6.0+) in the last 100 years .....	85
Figure 4-6 Histogram showing distribution of number of earthquakes .....	87
Figure 4-7 All earthquakes that could affect the Canterbury region.....	89
Figure 4-8 Filter applied in the EXSIM simulation .....	92
Figure 4-9 PLUM confirmation and warning .....	94
Figure 4-10 Population of the selected region in Canterbury .....	104
Figure 4-11 Failure Scenario A: removing one-third of the stations .....	106

Figure 4-12 Failure Scenario B: removing stations in a band .....	107
Figure 4-13 Fraction of shaking events with Correct Timely Alerts.....	109
Figure 4-14 Expected values of $CTA_{norm,g}$ .....	110
Figure 4-15 Expected values of False Alerts vs All Alerts.....	111
Figure 4-16 Expected value of Precautionary Alerts Vs All Alerts .....	112
Figure 4-17 Trade-off between expected $CTA_{norm,g}$ and expected $FA_{norm,g}$ .....	114
Figure 4-18 Values of Precautionary Alerts Vs All Alerts .....	115
Figure 4-19 Expected mean warning times for correct alerts .....	118
Figure 4-20 Distribution of WTCTA: box plots of warning times of all correct <b>timely</b> alerts .....	119
Figure 4-21 Distribution of WTCA: box plots of warning times of all correct alerts .....	120
Figure 4-22 Population-weighted $CTA_{norm,g}$ PLUM warning radius=50 km .....	123
Figure 4-23 Population-weighted $CTA_{norm,g}$ , other PLUM warning radii .....	124
Figure 4-24 Population-weighted $FA_{norm,g}$ PLUM radius =50 km.....	125
Figure 4-25 Population-weighted $FA_{norm,g}$ , PLUM radius 30, 40, 60 km.....	126
Figure 4-26 Population-weighted $PA_{norm,g}$ .....	127
Figure 4-27 Population-weighted expected mean warning times for all correct alerts .....	130
Figure 4-28 Population-weighted expected mean warning times for all correct alerts .....	131
Figure 4-29 $CTA_{norm,g}$ for warning radius 50 km, with 33% of stations removed .....	133
Figure 4-30 $FA_{norm,g}$ for warning radius 50 km, with 33% of the stations removed .....	133
Figure 4-31 $PA_{norm,g}$ for warning radius 50 km, with 33% of the stations removed .....	134
Figure 4-32 Expected $WTCA_g$ with 33% of stations removed .....	136
Figure 4-33 Distribution of WTCA with 33% of stations removed.....	137
Figure 4-34 Distributions of WTCTA with 33% of stations removed .....	138
Figure 4-35 $CTA_{norm,g}$ with stations in band removed.....	141
Figure 4-36 $FA_{norm,g}$ with stations in band removed for warning radius 50 km .....	141
Figure 4-37 $PA_{norm,g}$ with stations in band removed for warning radius 50 km .....	142
Figure 4-38 Expected $WTCA_g$ with stations in band removed.....	145
Figure 4-39 Distributions of WTCA with band of stations removed.....	146
Figure 4-40 Distributions of WTCTA with band of stations removed .....	147
Figure 5-1 The selected Wellington region with $0.1^\circ \times 0.1^\circ$ grid .....	158
Figure 5-2 Wellington region with sensors .....	159
Figure 5-3 VS30 site classification values across the Wellington region .....	161
Figure 5-4 Distribution of population in Greater Wellington region .....	162
Figure 5-5 Computed grid location site classes.....	163
Figure 5-6 All 1306 earthquakes that could affect the Wellington region .....	165
Figure 5-7 Histogram showing distribution of number of earthquakes .....	166
Figure 5-8 PLUM algorithm with confirmation from confirming station to alerting station. ....	167
Figure 5-9 Expected values of $CTA_{norm,g}$ .....	172
Figure 5-10 Expected values of $FA_{norm,g}$ .....	174
Figure 5-11 Expected values of $PA_{norm,g}$ .....	175
Figure 5-12 Expected mean warning times for correct alerts .....	178
Figure 5-13 Distribution of WTCTA .....	179
Figure 5-14 Distribution of WTCA .....	179

Figure 5-15 Population-weighted expected values of $CTA_{norm,g}$ .....	182
Figure 5-16 Population-weighted expected values of $FA_{norm,g}$ .....	183
Figure 5-17 Population-weighted expected values of $PA_{norm,g}$ .....	184
Figure 5-18 Population-weighted expected mean warning times .....	186
Figure 5-19 Centroids of all earthquakes causing strong shaking in south Wellington City .....	188

# List of Tables

Table 3-1 Data latency - outlier cases (latency > 6s) .....	36
Table 3-2 Three stations selected for the Kaikōura earthquake, .....	40
Table 3-3 Three stations selected for the Lake Grassmere earthquake .....	41
Table 3-4 Parameters of fault model .....	44
Table 3-5 EXSIM parameters for simulation.....	56
Table 3-6 PLUM performance over whole Canterbury region, basic metrics.....	61
Table 3-7 PLUM performance over whole Canterbury region, normalised metrics.....	62
Table 3-8 Sensor station density in the Canterbury region. ....	63
Table 3-9 PLUM performance before and after adding 5 new stations: total CTA.....	66
Table 3-10 PLUM performance before and after adding 5 new stations: total FA.....	66
Table 3-11 PLUM performance when adding one new station at a location .....	67
Table 3-12 Population-weighted performance measures.....	71
Table 4-1 Conversion of $I_{A:MMI}$ and $I_{V:MMI}$ to $I_{MMI}$ .....	78
Table 4-2 Possible outcomes of PLUM.....	95
Table 4-3 PLUM parameters for Canterbury catalogue earthquakes experiment. ....	98
Table 4-4 Basic Performance Indicators for EEW performance at a single grid cell. ....	100
Table 4-5 Expected number of each PLUM outcome across whole region .....	116
Table 4-6 Expected $CTA_{norm}$ , $FA_{norm}$ and $PA_{norm}$ across whole region .....	117
Table 4-7 Mean $\pm$ standard deviation of distributions of $WTCTA$ .....	120
Table 4-8 Unweighted and Population-weighted $CTA_{norm}$ , $FA_{norm}$ , $PA_{norm}$ .....	128
Table 4-9 Unweighted and Population-weighted mean warning times ( $\pm$ standard deviation) .....	132
Table 4-10 Expected $CTA_{norm}$ , $FA_{norm}$ and $PA_{norm}$ with 33% of the stations removed.....	135
Table 4-11 Expected number of each PLUM outcome with 33% of the stations removed .....	135
Table 4-12 Population-weighted $CTA_{norm}$ , $FA_{norm}$ , $PA_{norm}$ with 33% of the stations removed .....	136
Table 4-13 Mean $\pm$ standard deviation of distributions of $WTCTA$ and $WTCA$ .....	138
Table 4-14 Expected $CTA_{norm}$ , $FA_{norm}$ and $PA_{norm}$ , with station in band removed .....	143
Table 4-15 Expected number of each PLUM outcome with station in band removed.....	143
Table 4-16 Population-weighted $CTA_{norm}$ , $FA_{norm}$ , $PA_{norm}$ with station in band removed .....	144
Table 4-17 Mean $\pm$ standard deviation of distributions of $WTCTA$ and $WTCA$ .....	148
Table 5-1 Expected number of each PLUM outcome .....	176
Table 5-2 Expected $CTA_{norm}$ , $FA_{norm}$ and $PA_{norm}$ .....	177
Table 5-3 Mean $\pm$ standard deviation of distributions of $WTCTA$ and $WTCA$ .....	180
Table 5-4 Population-weighted $CTA_{norm}$ , $FA_{norm}$ , $PA_{norm}$ .....	185
Table 5-5 Population-weighted mean warning times .....	187

# Chapter 1 Introduction

Aotearoa New Zealand sits on the boundary of the Pacific and Australian tectonic plates, making it part of the Pacific 'Ring of Fire', a region known for its high volcanic and seismic activity. It is vulnerable to frequent earthquakes. The recent 2011 Christchurch earthquake caused widespread damage and killed 185 people (Potter et al., 2015). It is identified as New Zealand's fifth-deadliest natural disaster (Brake, 2018). The 1931 Hawke's Bay earthquake and 1820's Southland tsunami are identified as the second and third deadliest disasters, respectively (Hill, 2016; Worldatlas, 2021).

The Canterbury earthquake sequence (Townend, Villamor, & Quigley, 2012; Quigley et al., 2016), which consists of the Darfield earthquake (Mw 7.1) in September 2010, the Christchurch earthquake (Mw 6.2) in February 2011, and many aftershocks, is the largest insured event in New Zealand. Additionally, it was the fourth most expensive insured global natural disaster at the time (Insurance Council of New Zealand, n.d.). The earthquake sequence consumed more than 15% of New Zealand's GDP (Berryman, 2012). The 2016 Kaikōura earthquake (Mw 7.8) caused major damage to the Kaikōura coast and also damaged buildings in Wellington (Kaiser et al., 2017a; Cubrinovski et al., 2020).

Although earthquakes are not yet predictable, rapidly detecting and analysing the earthquake ground shaking and issuing an alert to regions before the shaking has reached them is one way of mitigating some of the risks of earthquakes (Satriano, Wu, et al., 2011). Such alerts, even when there are only a few seconds of warning time, can enable people to take precautionary actions that will reduce injury and fatalities, and may enable automated actions in buildings and other infrastructure that reduce damage (Strauss & Allen, 2016; Papadopoulos et al., 2023). Currently, earthquake early warning (EEW) systems have been implemented in several countries (Allen & Melgar, 2019), including Japan (Kodera et al., 2021), Taiwan (Wu et al., 2021), South Korea (Sheen et al., 2017), Mexico (Suárez, 2022) and the US West Coast ((Lux et al., 2024).

In New Zealand, GeoNet (the geological hazard monitoring system of New Zealand) issues an earthquake notification within a few minutes of an earthquake's occurrence based on an automatic and very preliminary determined location, magnitude and depth. GeoNet delivers

earthquake information to the public via a mobile application and a website (GNS Science, 2024b). However, New Zealand has no official public earthquake early warning (EEW) system (Vinnell et al., 2023). Therefore, New Zealanders are interested in a nationally integrated earthquake early warning system (Becker et al., 2020a).

Currently, in New Zealand, a few EEW systems are in operation or testing, with limited access, including the Android Earthquake Alert System (AEA). Google deployed an experimental AEA system in New Zealand in April 2021 and it is providing alerts to Android users only (Vinnell et al., 2023). The AEA system relies on accelerometers in Android phones to detect earthquakes (Vinnell et al., 2023) and operates independently of GeoNet or New Zealand government systems (Tan, Leonard, & Johnston, 2021). Tan et al. (2021) emphasise that the AEA should not be confused with alerts issued by New Zealand civil defence authorities and highlight the need for public education on EEW alerts.

Prasanna et al. (2022) introduced a low-cost EEW system using a Raspberry Shake sensor network and deployed a citizen science network experimentally in Wellington, New Zealand. This cost-effective solution addresses the economic challenge of limited high-end sensor density in New Zealand. Chandrakumar et al. (2024, 2025) explore advancements in low-cost EEW systems using that network. The 2024 study focuses on integrating node-level processing with the PLUM algorithm while the 2025 study evaluates the performance of P-wave detection algorithms for that community-based EEW system. Successful EEW systems, such as the ones used in Japan and Southern California, use different EEW algorithms and strategies (Allen & Melgar, 2019). Different algorithms have different strengths and weaknesses (Cremen & Galasso, 2020). Each algorithm depends on particular properties of the sensor networks, so an algorithm that works well in one country, for example Japan, with its very dense network of sensors (Kodera et al., 2021), may not work as well in other countries with different sensor networks and properties. Therefore, an EEW system cannot be simply copied from one country to another. This thesis explores and evaluates one of the EEW algorithms —PLUM (Propagation of Local Undamped Motion) (Kodera et al., 2021), — in the New Zealand context to evaluate its strengths and limitations if it were used as part of a New Zealand EEW system.

An important component of the NZ context is the GeoNet sensor network, so understanding the network and identifying its limitations is a critical part of this research. The thesis will explore some relevant features of this network. Evaluating an EEW algorithm requires earthquake data. Recorded data of past earthquakes are helpful, but are also very limited. Using synthetic data allows a more comprehensive analysis that considers likely earthquake risk for the near-term future. The thesis will use a large catalogue of synthetic NZ earthquake data (Shaw et al., 2022) for the evaluation and discuss how the earthquakes are simulated for this research. The thesis will also identify a set of metrics for evaluating an EEW algorithm that would allow stakeholders such as the National Emergency Management Agency (NEMA) and the Natural Hazards Commission to make informed decisions whether or not to use the PLUM algorithm in NZ. For example, to evaluate any warnings, we need to decide how to determine whether a warning was on time or not, whether it was a correct alert or a false alarm, and whether there was a missing alert for an event. Finally, the thesis presents the effectiveness of the PLUM EEW algorithm for the Canterbury and Wellington regions for possible future events using the generated synthetic seismograms. This is a novel approach as this evaluates an EEW algorithm using a physics-based synthetic earthquake catalogue, which will be useful in designing EEW systems for locations that lack enough historical records.

This research is part of the Resilience to Nature's Challenges (RNC2): Earthquake and Tsunami Programme, funded by the Ministry of Business, Innovation and Employment, New Zealand. The programme's aim is to improve the earthquake and tsunami resilience of New Zealand communities through innovative earthquake science.

Before starting the work for this thesis, the author worked on the low-cost EEW project mentioned above with A/Prof Raj Prasanna at Massey University and Dr Caroline Holden at SeismoCity/VUW, exploring low-cost Earthquake Early Warning systems. The work the author did for that project was published in Prasanna et al. (2022), but is not part of this thesis.

## **1.1 Research goals and objectives**

The thesis addresses the following research goals:

1. Develop and demonstrate a methodology for evaluating earthquake early warning (EEW) performance under a wide range of possible future earthquake scenarios in New Zealand using a synthetic earthquake catalogue.
2. Evaluate the performance of an EEW algorithm in the New Zealand context, and identify its strengths and limitations as a component of a public EEW system for New Zealand based on the current GeoNet sensor network

In order to achieve these goals, the thesis addressed the following objectives:

1. Research existing EEW algorithms and their suitability for New Zealand and choose an appropriate algorithm for evaluating EEW performance using a synthetic earthquake catalogue.
2. Determine and evaluate the current GeoNet sensor network capabilities and limitations (e.g., communication latency, and sensor distribution) in relation to an EEW system.
3. Check the appropriateness of, and find appropriate parameters for, an earthquake ground motion simulator for the NZ environment.
4. Use the simulator to generate synthetic seismograms for a large number of potential earthquakes for at least two regions of New Zealand.
5. Quantify the effectiveness of the chosen EEW algorithm using appropriate performance metrics in the New Zealand context.

## **1.2 Organisation of the thesis**

The rest of the thesis is organised as follows: Chapter 2 describes methods used in state-of-the-art EEW systems, the EXSIM earthquake simulator and the New Zealand synthetic earthquake catalogue. Chapter 3 presents several preparatory experiments and analyses to address objectives 1, 2 and 3. Chapters 4 and 5 present an evaluation of the PLUM algorithm in the Canterbury region and in the Wellington region to address objectives 4 and 5. Chapter 6 concludes the thesis with a summary of the results of the analyses, recommendations for a future EEW system, and future work building on the results of the thesis.

Note: When the thesis uses the word “we”, it is referring to the work of the author, with input and suggestions from the supervisors.

# Chapter 2 Background and Literature Review

This chapter gives an overview of the existing body of research relevant to the proposed thesis. It discusses state-of-the-art international earthquake early warning systems and algorithms and the New Zealand context for EEW and outlines the EXSIM ground motion simulator and the New Zealand (RNC2) synthetic earthquake catalogue.

## 2.1 New Zealand context for an EEW system

New Zealand's geological, geographical, technological, and socio-economic factors offer a complex and unique environment for EEW. It is located along the edge of two tectonic plates Australian and Pacific, and experiences a wide range of seismic activities including shallow crustal earthquakes, deep subduction earthquakes and interface earthquakes in subduction zones (Gerstenberger et al., 2024). The two long and narrow main islands, North and South, separated by Cook Strait provides additional uniqueness including the ways seismic waves can travel and how alerts are best distributed. High geological variability (Ballance, 2017) including diverse ground conditions (Perrin et al., 2015) should be considered when tailoring an EEW system for NZ. Different soil types, faults, and regional seismic characteristics can affect seismic wave propagation (Kaiser et al., 2024) and can influence both detection and warning accuracy and effectiveness. Population distribution is another key factor as the population of NZ is concentrated in a few big cities, especially Auckland, Wellington and Christchurch and each city has a different geological environment. While New Zealand has good telecommunication coverage in urban areas, rural and remote regions are less well served, which is another challenge for EEW effectiveness. The existing GeoNet sensor network provides the foundation for an EEW system; however, the high cost of expanding sensor density would be a challenge in developing a network as dense as Japan's where an EEW system has been successfully implemented. It would be difficult to justify the investment at a similar level of sensor network as Japan due to New Zealand's small population and geographical spread compared to Japan, given the country's economic constraints.

Additionally, social factors should be considered while developing an EEW system, which will be different in New Zealand from other countries. These unique characteristics necessitate consideration for EEW implementation to balance technological advancements with the country's specific seismic, geographic and socio-economic challenges.

The two organisations that would be key to a national EEW system for New Zealand are the National Emergency Management Agency (NEMA) and GNS Science, which runs the GeoNet project. Other important organisations include the Civil Defence Emergency Management Groups (CDEM), and other emergency service providers (such as Fire and Emergency New Zealand and NZ Police).

According to the Civil Defence Emergency Management Act 2002 (CDEM, 2002) and National Civil Defence Emergency Management Plan 2015 (Mateparae, 2015), NEMA has the authority to issue earthquake or tsunami warnings and advisories via the national warning system. The national warning system is used to disseminate all hazard situations in New Zealand and Figure 2-1 illustrates its structure for earthquake and tsunami events (NEMA, 2020a; 2020b). The New Zealand Emergency Mobile Alert service (Stuart-Black, 2021) is used to disseminate national warning messages as appropriate (NEMA, 2021).

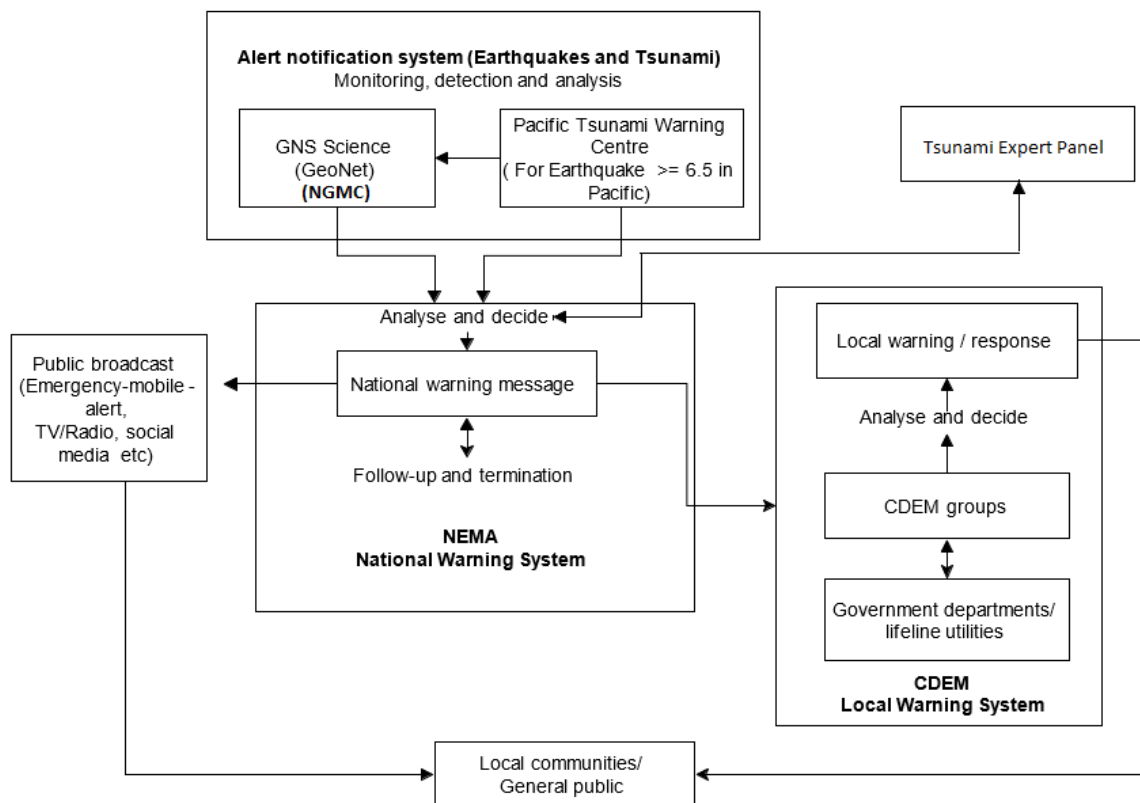


Figure 2-1 The National Warning System for Earthquakes and Tsunami

GNS Science operates the GeoNet programme, which manages the New Zealand geo-hazard sensor network (Petersen et al., 2011), and GNS Science is the lead monitoring and detection agency for earthquake information. Just after any earthquake event, GNS Science, via the GeoNet project, provides earthquake monitoring information to subscribers and the general public using the GeoNet website, smartphone apps, emails, and an RSS (a protocol for delivering regular updates to subscribers) feed. According to the National Civil Defence Emergency Management Plan, GNS Science must report an earthquake event to NEMA within five minutes of the event. After analysing the report, NEMA may decide to issue an advisory through the national warning system. Currently, GeoNet does not provide any early warning for earthquakes, and information is provided to stakeholders to help them manage post-event situations (Mateparae, 2015).

Future early warning systems should facilitate the requirements of a range of end-users, including rapid responders, industry, national agencies and the general public. Becker et al. (2020a) presented a sectoral analysis of the perceived benefits and challenges of a New Zealand EEW system. According to that research, requirements differ between sectors, but a

national earthquake early warning system could be beneficial to the general public of New Zealand. According to findings from a survey of 3,084 people (Becker et al., 2020b), 97% believed an EEW system would be useful or somewhat useful. 51% of them indicated the preferred threshold level to receive an alert as moderate shaking and 25% wished to receive an alert if the shaking intensity is strong or above. However, other end-users may have different requirements specific to their role, such as warnings for automated actions in large buildings, industrial plants, or major infrastructure. For example, both the US West Coast ShakeAlert System (Kohler et al., 2017) and the Japan EEW system (Yamasaki, 2012) use two types of alerts: one to the general public and one to expert users. Accordingly, the general public will receive the alert with limited essential information and only for significant events. Institutes, rapid responders, and industry will also obtain customised alerts. Becker et al. (2020b), found that requirements were also related to their geographical locations and experiences of past earthquakes. For example: varying threshold levels to warn were suggested by the emergency management sector such as MMI 5-7 for Napier; MMI 6 for Auckland; MMI 4-5 for Wellington; MMI 4-6 for Christchurch.

Google Android Earthquake Alert has been issuing alerts to Android mobile phone users of New Zealand since April 2021. However, the Google AEA is not endorsed by NEMA. Tan et al. (2023) conducted a study to understand whether the New Zealand public would continue to use the Android EEW. This study highlighted potential confusion among users about the source of the alerts and raised questions about responsibility and liability for the EEW system. There is very little information available about the Google algorithm or the performance of the NZ Android EEW.

Vinnell et al. (2023) surveyed 3,150 New Zealanders after two widely received AEA for Magnitude 5.3 and 5.9 earthquakes in October 2021 to assess public knowledge, perceptions, and responses. The vast majority of participants found the alert useful or somewhat useful, consistent with Becker et al. (2020b) previous survey findings. Vinnell et al. (2023) highlight the need to tailor EEW to New Zealand's unique seismic profile, population density, and prior earthquake experience. They mention that setting the threshold too low could lead to warning fatigue, reducing public responsiveness. The information provided in alerts should also be relevant to New Zealand. For example, the Google AEA issued to NZ includes safety tips that may not be relevant, such as prioritizing cleaning up spilled medicine over moving to

higher ground in case of a tsunami. This indicates the importance of adapting EEW systems to local contexts.

In addition to us there are two other research groups working on earthquake early warning algorithms in New Zealand. The GNS Science Rapid Characterisation of Earthquake and Tsunami (RCET) programme group is exploring the adaptation of the FinDer algorithm as a tool for earthquake response in New Zealand (Andrews et al., 2024). The CRISiSLab group at Massey University is exploring the use of low-cost sensors (e.g., Raspberry shake) to issue EEW in New Zealand and testing the PLUM algorithm on these sensors (Prasanna et al., 2022; Chandrakumar et al. (2024); Chandrakumar et al. (2025)).

## **2.2 Earthquake early warning algorithms**

The available earthquake early warning algorithms can be classified into four main types: point source, finite fault, ground motion and on-site algorithms.

### **2.2.1 Point source algorithms**

The earthquake source is represented as a simple point in space for point source algorithms. The algorithms use a few seconds of P-wave data from a small number of stations near the epicentre to determine an earthquake location, time, and magnitude (Allen & Melgar, 2019). Point source methods are among the fastest EEW techniques because they use the earliest seismic arrivals (P-waves) and require only short data windows. The four main steps of the point source method are (Satriano, Wu, et al., 2011):

- detecting an earthquake event and determining the location.
- estimating the magnitude of the event.
- predicting the ground motion for the target sites.
- delivering the alert.

The algorithms require data from at least two stations, and the accuracy of location and magnitude estimation is improved when data comes from more stations (Allen & Kanamori, 2003). These algorithms use 0.5 to 5 seconds of P-wave data (Allen & Melgar, 2019). Satriano, Wu, et al. (2011) state that magnitude estimation is more complex than event detection and localisation. After determining the location and magnitude, ground motion prediction

equations (GMPE) are used to predict the shaking intensities for the target area. Then the system can deliver an alert to the interested parties. However, the final outcome will depend on the ambiguities of the characterisation of location, magnitude and peak ground motion prediction (Satriano, Wu et al., 2011). A limitation of this approach is that the algorithm will be saturated for large earthquakes, typically magnitude greater than 7 (Kanamori, 2005; Zollo et al., 2006). For an example, the 2011 M 9.0 Tohoku-Oki earthquake was underestimated by point source EEW (Colombelli et al., 2012). Such saturation is likely caused by relying on only a few seconds (up to 5 seconds) of the P-wave, which is insufficient to capture the full rupture process of a large earthquake (Rydelek et al., 2007; Festa et al., 2008). Colombelli et al. (2012) suggested progressively expanding the P-wave time window as a solution to this saturation issue. However, Trugman et al. (2019) indicates that it is challenging to detect the full rupture process and respective magnitude in real-time. Therefore, using predictions from a saturation model, Trugman et al. (2019) developed a Bayesian framework to estimate posterior uncertainties in real-time magnitude assessments. However, ground motion prediction using a point-source rather than an extended rupture is still challenging. The relationship between P wave peak displacement and earthquake magnitude was explored by Trugman et al. (2019) using a large dataset of M4.5–9.0 earthquakes recorded by the Japanese strong-motion networks. The time evolution of P wave peak displacement in this dataset shows a clear pattern of initial growth, which does not align with deterministic models of earthquake rupture.

There are several point source algorithms available in the literature, including Earthquake Alert Systems (ElarmS), (Allen & Kanamori, 2003), Earthquake Point-Source Integrated Code (EPIC) (Chung, Henson, & Allen, 2019) and Probabilistic and Evolutionary Early Warning (PRESTo) algorithms (Satriano, Elia, et al., 2011).

### **ElarmS algorithm**

The ElarmS algorithm needs data from three or more stations to deduce the location and magnitude of the earthquake. It uses a short-term-average/long-term-average (STA/LTA) algorithm, which needs 4 seconds of P wave data to estimate the magnitude properly for events higher than magnitude 4.5. Several versions of the algorithm have been tested in California, Oregon, Washington, the Pacific Northwest of the United States and South Korea

(Allen et al., 2009; Allen & Melgar, 2019). The algorithm was also tested using Japan, Taiwan, Italy, and Alaska seismic data. In Southern California, where the station spacing is about 50 km, ElarmS is capable of issuing a warning within 30 s (average 14.9 s) depending on the threshold (Wurman, Allen, & Lombard, 2007).

ElarmS-2 is an improved version of the algorithm that includes new features: a new waveform processor module and a feature preventing false alerts due to teleseismic events (Serdar Kuyuk et al., 2014).

It has been further improved to ElarmS-3 and it was used by the ShakeAlert EEW system of the U.S. West Coast. ElarmS-3 is capable of preventing false alerts from other ElarmS algorithms by introducing new trigger filters, such as an amplitude-based check and a horizontal-to-vertical ratio check (Chung et al., 2019).

### **EPIC algorithm**

The EPIC algorithm has been developed by further improving the ElarmS-3 by introducing a filtering mechanism to suit different data sampling rates (Chung et al., 2019; Kohler et al., 2020).

ElarmS-3 uses the same filter coefficient for all sampling frequencies. EPIC's filter coefficients are based on the observed frequencies and remove the average amplitude calculated using the previous 60 seconds of the data from each sample. According to the test results for the ShakeAlert system, EPIC has the capability of producing an alert with 3, 6 and 9 seconds warning times for locations 20, 30, and 40 km from an epicentre, respectively (Chung et al., 2019). However, sparse station coverage can challenge ShakeAlert, as the EPIC algorithm requires P-wave triggers from at least 40% of stations within the radius of the farthest trigger to issue an alert, with a minimum of four triggering stations (Lux et al., 2024).

### **PRESTo algorithm**

PRobabilistic and Evolutionary early warning SysTem (PRESTo) uses a probabilistic approach to estimate earthquake location and magnitude from three or more stations. It can determine the approximate location and magnitude within five to six seconds. The solution evolves and stabilises within 10 seconds. RTLoc (Satriano, Lomax, & Zollo, 2008) and RTMag (Lancieri &

Zollo, 2008) techniques are used to determine location and magnitude within PRESTo algorithm (Satriano, Elia et al., 2011).

### **2.2.2 Finite fault algorithms**

Earthquakes are not actually point sources as they involve slip over a fault area that increases in size along with magnitude. Thus, traditional point source EEW algorithms do not work well for large earthquakes ( $M > 7$ ) due to the complex and extended nature of fault rupture (Ruhl et al., 2017; Kodera et al., 2021). To address this limitation, finite fault algorithms have been developed to generate EEW by incorporating earthquake location and fault characterisation by modelling the spatial and temporal evolution of fault rupture. These algorithms are not saturated for large earthquakes ( $> M7$ ). However, finite fault algorithms take more time to issue an alert than point source methods (Ruhl et al., 2019; Allen et al., 2019).

Unlike point source approaches, finite fault algorithms track the developing rupture in real-time. They produce a time series of alerts that progressively refine estimates of magnitude and fault dimensions. While preliminary alerts may be issued quickly, the most accurate and stable solutions can only be delivered once the rupture has substantially or fully completed.

In addition to EEW, these algorithms are often used as rapid response tools. Rapid response tools are designed to quickly assess and communicate the potential impact of an earthquake immediately after it occurs, providing critical information for emergency response and recovery efforts.

Finite fault algorithms use both seismic and geodetic approaches.

Seismic-based finite-fault algorithms are designed to overcome the limitations of point source methods by modelling the spatial extent of rupture using data from dense seismic networks. These algorithms extract fault geometry by analysing ground motion observations, slip patterns, rupture directivity, and the evolving shape of the seismic wavefield. For instance, FinDer algorithm (Böse, Heaton, & Hauksson, 2012) uses peak ground acceleration (PGA) observations from ground motion sensors and compares them with pre-computed templates to determine rupture length, position, and orientation. Another example is the iterative deconvolution and stacking (IDS) approach, which automatically inverts strong-motion records to rapidly estimate the finite-fault models (Zheng et al., 2020).

Geodetic approaches use Global Navigation Satellite System (GNSS) data to complement traditional seismic methods. GNSS sensors provide direct measurements of static ground motion, making them particularly effective for estimating fault length and magnitude during large earthquakes (Ruhl et al., 2017). This is especially valuable for events exceeding magnitude 7, where seismic methods may saturate (Crowell, 2024). However, the relatively high noise floor of GNSS sensors limits their sensitivity to smaller magnitude earthquakes, reducing their effectiveness for events below approximately magnitude 6.5 (Ruhl et al., 2019).

The following discusses several finite fault algorithms including FinDer (Böse, Heaton, & Hauksson, 2012), G-larmS (Grapenthin, Johanson, & Allen, 2014a), G-FAST (Crowell et al., 2018) BEFORES (Minson et al., 2014), and REGARD algorithms (Kawamoto et al., 2017).

### **FinDer algorithm**

The Real-time Finite Fault Rupture Detector (FinDer) algorithm uses an image processing approach to estimate fault rupture extent for large earthquakes (Böse et al., 2012, 2018, 2023). It uses ground motion data from seismic sensors and compares the data with predetermined templates. A correlation-based matching technique determines the current centroid position, length, and strike. It can provide EEW and shake-maps by updating source geometry estimates continuously. First, it detects maximum ground motion amplitudes and applies site corrections. Then it performs map projection and spatial interpolation of amplitudes. After the near and far source classification, it compares the map with a set of pre-calculated templates. Then, a direct pattern search algorithm is used to minimise the misfit. . FinDer has been successfully integrated in the ShakeAlert EEW suit since 2018 along with the EPIC point source algorithm for the west coast of the United States (Kohler et al., 2020) and has also been implemented for Switzerland (Massin, Clinton, and Böse, 2021) and Costa Rica (Porrás et al., 2021). Andrews et al., (2024) tested FinDer algorithm for New Zealand. It works well for Christchurch, where station density is high, and it computed the first solution within 7 seconds of the origin time.

### **G-larmS algorithm**

The Geodetic Alarm System (G-larmS), a.k.a. rtGPS static offset inversion algorithm, determines earthquake magnitude and finite fault details using least-squares inversion of the

static offset of slip (Grapenthin et al., 2014a; Grapenthin, Johanson, & Allen, 2014b) using GNSS data. GNSS data are typically relatively noisy; therefore, G-larmS is designed to rely on triggers from more sensitive detection algorithms (Ruhl, 2017) such as the ElarmS point source algorithm.

It generates displacement time series from the geodetic data and then removes the outliers and excessive data gaps. It filters any noise and extracts the static offsets. Finally, it determines earthquake parameters using the inversion of these offsets. It can provide a distributed slip model and magnitude estimate within 14 to 24 seconds after an earthquake. It can process multiple events in parallel and might be useful during complex earthquake sequences.

### **G-FAST algorithm**

The Geodetic First Approximation of Size and Time (G-FAST) algorithm (Crowell et al., 2018) consists of a combination of different approaches. The processing starts after receiving a triggering signal from other seismic detection algorithms (e.g., ElarmS), which use seismometer data. Then, peak ground displacement (PGD) from Global Positioning System (GPS) displacement waveforms is used to determine magnitude and depth. After the fault orientation is obtained by using a geodetically derived focal mechanism, a discretised fault plane is created, and the static slip distribution is obtained by inverting it. It has been demonstrated that G-FAST can deliver magnitude estimates based on PGD within 22 seconds (first alert) of the earthquake's origin time and a stable solution within 40 to 60 seconds.

### **BEFORES algorithm**

The magnitude, spatial distribution of slip, and fault geometry can be captured in real-time using the Bayesian Evidence-based Fault Orientation and Real-time Earthquake Slip algorithm (BEFORES) (Minson et al., 2014). Bayesian inference is used to find the optimal fault geometry, and a simple analytical method is used to find the distributed slip models. The time taken to issue an EEW depends on the rupture duration and the duration of the wave propagation time from the source to the sensors. Real-time high-rate GPS data is used to solve this inversion problem using an analytical Bayesian approach. First, it considers the stations which detect P waves and estimate quasi-static offsets. Then it searches over strike and dip. Finally, it reports

magnitude ( $M_w$ ) and slip with uncertainties. The steps are repeated over time. It is assumed that the location is provided by the existing EEW method. BEFORES can provide results within 1 to 3 minutes after the earthquake's origin time, depending on the event's characteristics and the density of the GPS network and is therefore more useful for rapid response than EEW.

### **REGARD algorithm**

The Real-time GEONET Analysis system for Rapid Deformation (REGARD) algorithm (Kawamoto et al., 2017) uses a fault modelling approach based on the GNSS Earth Observation Network (GEONET) of Japan. The real-time GNSS positioning module updates the coordinate time series continuously and feeds it to the event detection module. If an earthquake occurs, the finite fault model is determined using the extracted co-seismic displacement field. Fault model estimation is updated continuously using a quasi-finite fault inversion routine. It can resolve a finite fault model for large earthquakes, avoiding instrumental saturation. It is capable of producing the first detection within 40-60 seconds and stable results within three minutes of the origin time..

### **2.2.3 Ground motion based algorithms**

Ground motion based algorithms predict seismic intensity for moderately distant sites using the observed local intensity (Allen & Melgar, 2019). They do not rely on an earthquake source model and are therefore less computationally expensive. The key advantages of these algorithms are that they are robust, rapid, and simple. However, they cannot determine detailed earthquake source parameters.

Both point source and finite fault algorithms first resolve the earthquake source parameters (i.e., location and magnitude) and then predict the ground motion to issue EEW. A problem with these algorithms is under-prediction and over-prediction (Hoshiaba, 2013; Hoshiaba & Aoki, 2015) and they find it difficult to handle simultaneous multiple events, especially during aftershock sequences. As a solution, Hoshiaba & Aoki (2015) proposed a ground motion numerical shake prediction method that does not need earthquake source characteristics.

## PLUM algorithm

One ground motion based algorithm is the Propagation of Local Undamped Motion (PLUM) algorithm (Kodera et al., 2018; Kodera, 2019). PLUM is derived from the numerical shake prediction method of Hoshiya & Aoki, (2015). PLUM is computationally inexpensive and easy to implement, compared to the numerical shake prediction method. It predicts shaking at distant locations (within a specific range) using ground motion observations from strong motion sensors. Figure 2-2 illustrates an example arrangement of sensors for PLUM within the specific region  $R_{reg}$  (the circular region with radius  $R$  centred at the observation station).

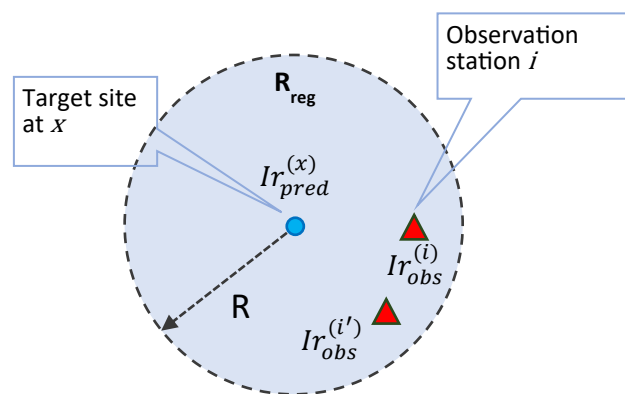


Figure 2-2 Plum algorithm (adapted from: Kodera et al., 2018)

PLUM predicts seismic intensity at the target point  $x$  using the observed intensity at location  $i$ . Empirical site amplification factors of both observation and prediction stations are taken into account. Seismic intensity at the predicted point is obtained by equation 2.1.

$$I_{pred}^x = \max_{i \in R_{reg}} \{I_{obs}^i - F_i^i\} + F_i^x \quad (2.1)$$

where:

$I_{pred}^x$  = predicted seismic intensity at  $x$

$I_{obs}^i$  = observed real-time seismic intensity of a station at  $i$

$F_i^i$  = site amplification factors at  $i$

$F_i^x$  = site amplification factors at  $x$

Japan's JMA EEW system currently uses PLUM and conventional point-source approaches to issue EEW. Each approach evaluates the ground shaking separately, and the maximum value is considered when issuing EEW (Kodera et al., 2020). The results show that PLUM produces

robust and accurate ground-motion predictions compared to the approaches which use source parameters and GMPE.

Performance assessment of the PLUM algorithm for Southern California was conducted by Cochran et al. (2019). They tested the algorithm using Southern California earthquakes that occurred from 2012 to 2017. They used a testing tool that was developed for the ShakeAlert EEW system. They found that PLUM can detect modified Mercalli intensity ( $I_{MMI}$ ) 4.0 and above earthquakes, which is an even lower threshold than the Japan PLUM EEW system's threshold of  $I_{MMI} \approx VII-VIII$  (Kodera et al., 2020). The  $I_{MMI}$  scale is used to find the impact of an earthquake, and the  $I_R$  represents the real-time seismic intensity. In Japan,  $I_R$  is used for PLUM, but  $I_{MMI}$  was used for this assessment by converting the ground motion to  $I_{MMI}$  according to the relationship proposed by Worden et al. (2012). Cochran et al. (2019) used  $I_{MMI}$  as it is the standard way to represent ground shaking in the United States despite the fact that  $I_{MMI}$  approach is more sensitive to noise spikes than the  $I_R$  approach. However, they modify PLUM to reduce false alerts induced by noise spikes. PLUM is modified by considering signals coming from two stations to trigger the algorithm, which helps to reduce false alerts that can be generated by noise on a single station. The test results demonstrate that all M 5.0+ earthquakes were detected when there are two stations or more within 60 km. Cochran et al. (2019) suggested examining different approaches and thresholds for scenarios different to southern California, such as different station densities, event characteristics and velocity models.

Kilb et al. (2021) reported the performance of PLUM for two West Coast (USA) data sets. It provided the same or better performance (mean detection time 8s) as other EEW methods for the 2012–2017 (California, Oregon, Washington) data set. However, the mean detection time was increased for the 1999–2015 earthquake data set where station density is lower than the previous data set. They further improved the algorithm by introducing two features to the work of Cochran et al. (2019). They used co-located velocity and acceleration sensors to improve the detection by reducing the noise and eliminating other issues of individual sensors. Another new feature is two-station event detection. A nearby sensor validates the initial triggering to issue an alert within 15 seconds. It is useful for reducing false alerts due to uncorrelated, time-separated noise on two far-away stations. If the station density changes, it is necessary to adjust the fifteen-second time window as suggested by Kilb et al. (2021).

Böse et al. (2022a) evaluated the PLUM performance for the foreshock and aftershock sequences alongside assessing the ShakeAlert EEW system's EPIC and FinDer algorithms, as the PLUM will be considered for the future ShakeAlert integration. They used synthesised earthquake sequences with two earthquakes occurring within 3 minutes of each other. Their finding showed that the PLUM performs well for near-simultaneous events and is effective in identifying pockets of ground motion. In this kind of complex situation, PLUM's ground motion estimates work better than the methods that rely on source parameter determination. Böse et al. (2022a) mentioned that PLUM works well for temporally and spatially close earthquakes, as they are difficult to identify and issue EEW using traditional point source methods.

Despite the advantages of the PLUM, there are several limitations. Since it does not determine the location of the earthquake, the maximum warning time is limited by the chosen warning radius, whereas source characterising approaches could produce longer warnings to distant locations once they locate the earthquake. If the warning radius is increased beyond the typical range, it may violate the PLUM assumption of undamped motion and result in a higher percentage of false alerts. Additionally, because PLUM detects earthquakes using secondary-waves (S-waves), it requires a higher ground motion threshold than primary-wave (P-wave) based algorithms, potentially resulting in delayed detection.

### **ALPHA algorithm**

Another ground-motion based algorithm is the Approximation by Local Pseudo-Hypocentre Attenuation (ALPHA) method (Kodera, 2019) (see Figure 2-3). The PLUM algorithm assumes that body waves propagate without attenuation. Therefore, PLUM can only predict seismic intensity for a short distance (i.e., less than 30 km). The ALPHA method was introduced to predict seismic intensities for longer distances by estimating the attenuation. The ALPHA method also works well for complex scenarios such as simultaneous earthquakes. It predicts seismic intensities for distant sites using GMPE by assuming that the point source hypocentre is located below the reference observation station. However, it is computationally expensive than the PLUM (Kodera, 2019), and according to our best knowledge, ALPHA is still not being used in any operational EEW.

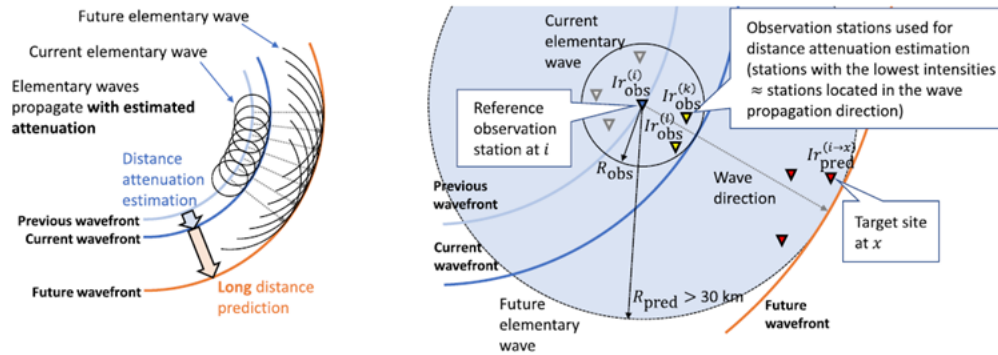


Figure 2-3 Alpha algorithm (source: Kodera 2019)

### 2.2.4 On-site algorithms

On-site algorithms use data from a local sensor to generate EEW for the same site (Tsunoi et al., 2024). They do not depend on other sensor data or communication infrastructure like the other algorithm types. Generally, this method uses preliminary P-wave data to predict S-wave amplitudes. It triggers an alert when the following S-wave or peak shaking is predicted to be above a certain threshold. It can issue an EEW using only one instrument. Nevertheless, it may generate more false alarms because it is difficult to eliminate noise or other issues using only a single device (Caruso et al., 2017). It may fail for large earthquakes and for complex aftershock sequences, since large earthquake rupture processes are complex, and it can be difficult to differentiate P and S waves properly.

The Taiwan National Centre for Research on Earthquake Engineering has developed an on-site algorithm that uses six P-wave features of the vertical component of the initial three seconds of data. These features include: predominant frequency, peak ground acceleration, peak ground velocity, peak ground displacement, cumulative absolute velocity and integration of the squared velocity. This data is used to predict the peak ground acceleration using a regression model based on a support vector machine (Hsu et al., 2016). National Taiwan University developed an on-site algorithm named P-alert, which also uses the first three seconds of P-wave data. It is able to estimate peak ground motion velocity and magnitude by computing a ground-motion period parameter and a vertical displacement amplitude parameter from the P-wave data (Wu & Kanamori, 2008). P-alert sensors have been installed in New Zealand by multiple organisations including Kiwi Rail, Callaghan Innovation, Wellington Water and Wellington Electricity (JENLOGIX, 2021). In New Zealand,

P-alert sensors are mainly used for structural health monitoring, meaning there is no EEW capability evaluation for New Zealand in the literature.

To address the limitations of on-site algorithms, Zollo et al. (2014) proposed an integrated approach that combines on-site algorithms with the point-source PRESTo algorithm. This hybrid method aims to enhance the robustness and overall performance of EEW systems by leveraging the strengths of both approaches.

### **2.2.5 Selection of an algorithm to test for NZ**

There are four main types of EEW algorithms, as described in previous Sections 2.1.1 to 2.1.4. As mentioned in Section 2.1.3, ground motion based algorithms have a number of advantages—for example, they are computationally inexpensive, fast, simple to implement, and capable of handling complex earthquake sequences.

While both PLUM and ALPHA algorithms are ground motion–based and do not rely on source parameters, they work quite differently. PLUM is quick and effective for nearby areas, using real-time shaking data from nearby sensors to issue alerts. But its range is limited by the warning radius due to its undamped motion assumption (which does not account for wave attenuation). ALPHA, on the other hand, estimates how shaking will decrease with distance by considering each station as a kind of pseudo-hypocentre. Thus, it could predict shaking farther away, which is especially useful for distant sites. ALPHA is more computationally demanding and isn't yet part of any operational EEW system. For the Japan 2011 M9 Tohoku-Oki earthquake scenario testing, the total computational times of ALPHA were 44 times those of PLUM (Kodera, 2019). PLUM and ALPHA could complement each other—PLUM gives fast local warnings, and ALPHA extends coverage to areas farther from the source. Therefore, Kodera (2019) suggested that using both could improve overall early warning performance.

Therefore, a ground motion based algorithm, named PLUM (see Section 2.1.3 for the details of the PLUM) has been selected for testing the GeoNet network early warning capabilities and testing ALPHA combining with PLUM will be an interesting future work.

The PLUM algorithm has several advantages, and consequently, the Japan Meteorological Agency (JMA) has been using it since 2018 in the Japan EEW system (Kodera et al., 2020). Based on its performance in the Japan EEW system, Kodera et al. (2020) suggested that the

PLUM method is more reliable and robust than point source algorithms in complex earthquake scenarios. It can also better capture high-stress drop earthquakes than other conventional algorithms (Kilb et al., 2021). It is capable of handling simultaneous earthquakes (Böse et al., 2022a), which is also useful during the aftershocks. Böse et al., (2022a) mentioned that the PLUM EEW algorithm is being examined for potential inclusion in the US West Coast ShakeAlert EEW system. Minson et al. (2020) performed real-time testing for PLUM in California, Oregon, and Washington. During the 2019 Ridgecrest earthquake sequence (M 6.4 and 7.1). They observed that the PLUM EEW is faster than the ShakeAlert system's first alerts. High data latency during the M7.1 earthquake had effect on the ShakeAlert performance, but not on the PLUM.

PLUM performed well in Japan with a high-density sensor network (Kodera et al., 2020). However, NZ does not have a high density of sensors like JMA. US data (Cochran et al., 2019, Kilb et al., 2021) shows that PLUM can work with a lower-density sensor network compared to Japan's network. Therefore, it is worthwhile to study the PLUM algorithm for New Zealand settings.

The PLUM algorithm was selected as a suitable starting point for exploring the evaluation of an EEW algorithm using a large synthetic earthquake catalogue for NZ. Its simplicity (being computationally inexpensive) makes it easier to implement and test compared to more complex algorithms. Additionally, PLUM has relatively modest data requirements as it does not rely on detailed source information and instead detects ground motion primarily using S-wave data. While the choice of algorithm is important, the methodology for evaluating EEW performance using a large synthetic catalogue is even more critical in alignment with the original objectives of the main RNC 2 Earthquake and Tsunami project. Internationally, Japan employs both point source approaches and the PLUM method within its EEW system. The US West Coast ShakeAlert EEW system integrates both the point source EPIC algorithm and Finite Fault FinDer algorithm. In New Zealand the FinDer finite fault algorithm has been tested by the R-CET project team at GNS Science, and currently, GeoNet issues earthquake information using earthquake source characterisation methods. Given these factors, PLUM offers a practical and effective foundation for testing EEW capabilities in the New Zealand context as part of this thesis.

## 2.3 Earthquake early warning systems

There are a few countries that have public EEW systems, including Japan, Taiwan, South Korea, Mexico, and the West Coast of the United States. Some other countries have restricted EEW systems that provide warnings only to specific organisations or user groups (Allen & Melgar, 2019).

The Japanese system uses hybrid methods to issue EEW using more than 4000 seismic monitoring stations (Tajima & Hayashida, 2018). During the Tohoku earthquake (M9.0) in 2011, the Japan EEW system successfully issued an early warning to the nearest region from the epicentre, before the arrival of the damaging S-waves. However, the system also issued many false warnings for aftershocks. Lack of data from seismometers in disaster areas (due to blackouts, communication failures, etc.) was one of the major causes of those false alarms (Nishimae et al., 2016). Until 2018, they used a point source approach to issue EEW. Since 2018, Japan's EEW system has been enhanced by introducing the PLUM algorithm to the existing EEW system to address some of the previous issues (Kodera et al., 2020).

The Taiwan Central Weather Administration (CWA) operates an EEW system which takes about 22 seconds to generate an EEW alert. Therefore, it can only provide EEW for areas at least 70 km from the epicentre (Wu & Kanamori, 2005; Xu et al., 2017). In addition to the CWA official EEW, Taiwan uses two other EEW systems based on the P-alert method (Wu et al., 2013; Xu et al., 2017). This onsite P-alert method can in theory provide quicker warnings than the CWA official EEW system. Accordingly, Chen et al. (2015) proposed incorporating low-cost P-alert sensors to the CWA EEW system. Taiwan has also recently implemented an Earthquake and Tsunami Submarine Cable Observation System in the Eastern Sea of Taiwan to assist the CWA to improve the earthquake and tsunami warning capabilities (Hsiao et al., 2014; NEC, 2021).

The Korean Earthquake Early Warning System (KEEWS) was established in 2015. The ElarmS algorithm (Wurman et al., 2007) was used by KEEWS initially and then improved by using ElarmS-2 algorithm (Serdar Kuyuk et al., 2014). According to the initial testing of the United States ElarmS algorithm for Korean data, the parameters should be adjusted for local crustal structure and regional seismicity. Therefore, they changed the magnitude scaling relationship and phase association criteria of the original ElarmS-2 which was set originally to the

Californian environment. It is proposed to increase its 188 seismic sensor stations to 300 to reduce the warning time (Sheen et al., 2017).

The seismic alert system of Mexico City (SAS) and the seismic alert system of Oaxaca city (SASO) were introduced in 1991 and 2003, respectively. In 2012, both were integrated into a single Mexican seismic early warning system known as SASMEX (Cuellar et al., 2014). A modified SLA/LTA method is used by them as the detection algorithm.

In California, the ShakeAlert EEW system has been tested for public alerting since October 2019. It uses both the EPIC and FinDer algorithms (Lux et al., 2024) complementing each other. The ShakeAlert system uses the US Wireless Emergency Alerts System (a method to deliver emergency alerts to mobile devices), a mobile app (MyShake) and Google Android OS to disseminate the EEW. The US West Coast ShakeAlert system started operation in May 2021, and has issued public EEW alerts to California, Oregon and Washington (USGS, 2021a). ShakeAlert was developed by the United States Geological Survey (USGS) in a collaboration of State and university partners. The Android earthquake alerts team of Google is one of the licensed-to-operate partners of the ShakeAlert product (USGS, 2021b), and Google does not issue their own EEW (i.e., AEA) in the US.

Google has been testing its own Android Earthquake Alerts System in Greece and New Zealand since April 2021 (Spooner, 2021), followed by testing in other countries. It uses accelerometer data coming from android phones. If shaking is detected, it sends the data to an earthquake detection server. The location and magnitude are determined using the data from the Android phones, and the alert is disseminated to other Android phones.

## **2.4 GeoNet sensor network and capabilities**

GeoNet is a geohazard monitoring network for NZ, which provides detection and monitoring for earthquakes, tsunamis and other geological hazards, and it is managed by GNS Science (Petersen et al., 2011). It captures geophysical data streams from sensors around the country (GNS Science, 2024c). The seismic data stream is analysed in near real-time by the 24/7 National Geohazard Monitoring Centre, and the data is also available to researchers. The GeoNet mobile application and the website provide geological hazards information to the

general public. They provide earthquake information, including location, magnitude and depth, usually within a few minutes of the earthquake occurrence (GNS Science, 2024b).

GeoNet runs several types of sensors including broadband seismometers, strong motion sensors, short-period seismometers, coastal sea level gauges, and geodetic sensors (GNSS/GPS), as shown in Figure 2-4 (GNS Science, 2024c). GeoNet maintains the New Zealand National Seismograph Network, Regional Seismograph Networks and National Strong Motion Network sites. Strong motion sensors have been selected to be utilised in this research because they can capture a broad range of frequencies and handle a large range of shaking amplitudes, including very strong shaking, without saturating. In September 2023 (when the sensor location data was downloaded to run the experiments in chapters 4 and 5), there were 344 strong motion stations around New Zealand. The work in this thesis focuses on the strong motion sensor network.

To determine station spacing across New Zealand's entire GeoNet strong-motion sensor network, we computed each station's distance to its nearest neighbour (344 stations). The distribution of the data is shown in Figure 2-5. Distances range from 0.2 km to 59.8 km, with a mean of 14.5 km (standard deviation  $\sigma = 13.8$  km) and a median of 10.9 km (interquartile range 1.7 km–25.0 km). The long upper tail reflects sparse station placement in remote regions, while most stations are within  $\sim 11$  km of another, indicating denser coverage in populated areas.

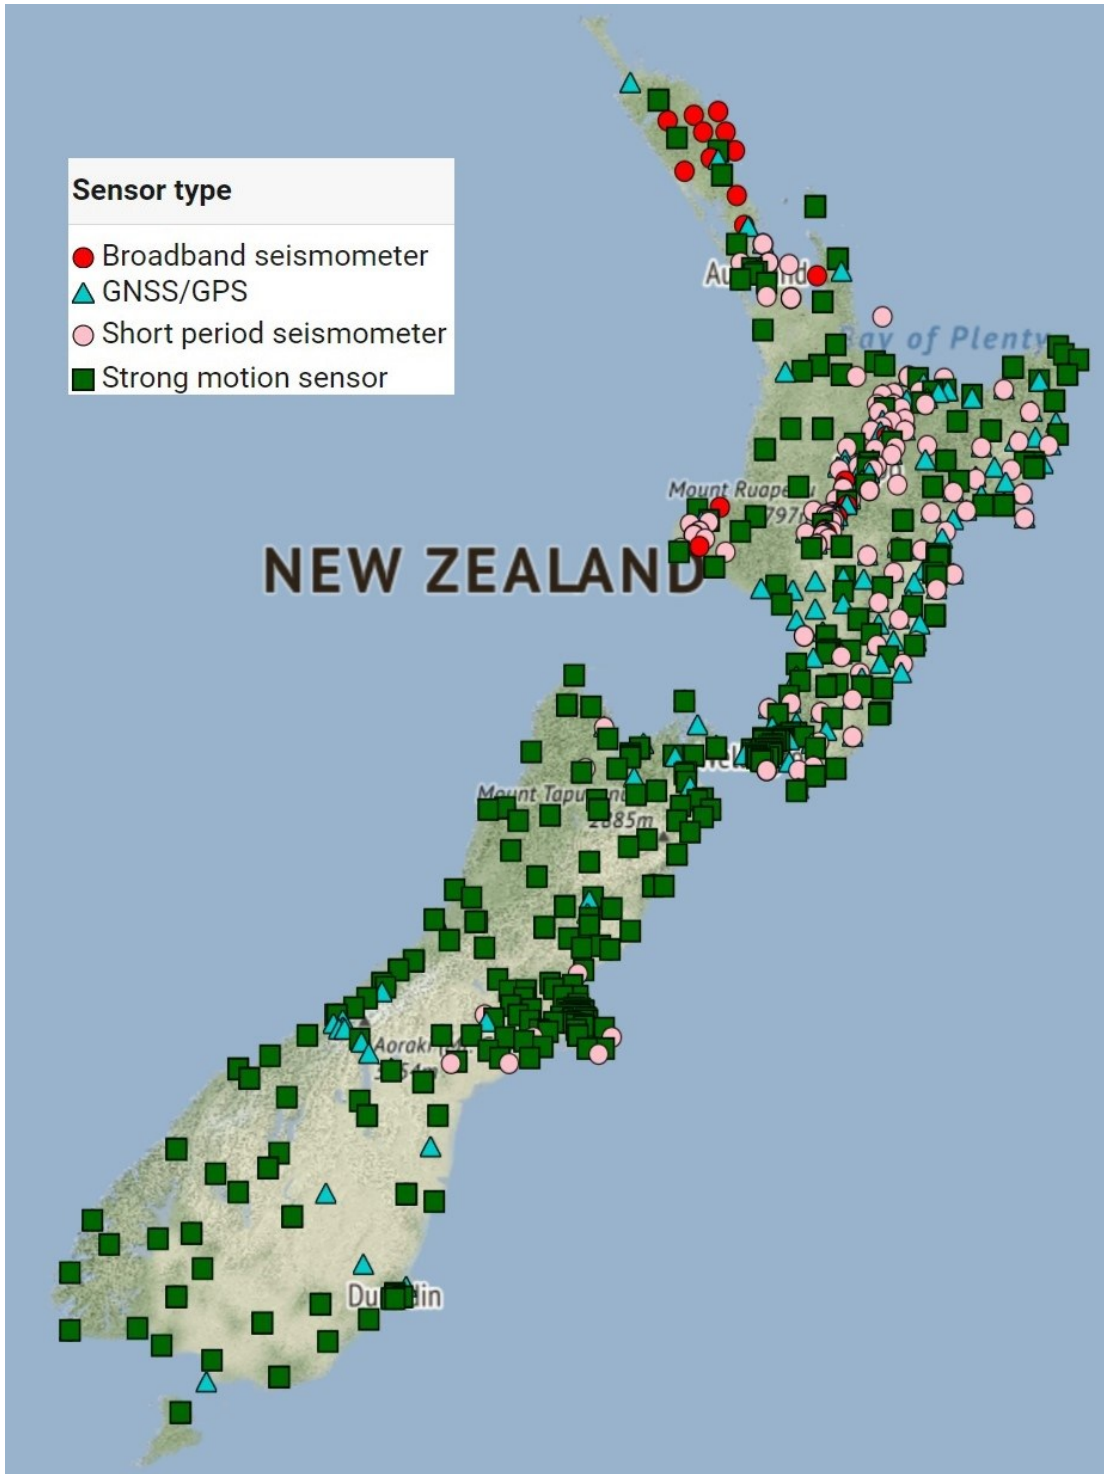



Figure 2-4 GeoNet seismic and GPS sensor map (North and South Islands), source GNS Science (2024c), <https://www.geonet.org.nz/data/network/sensor/search>



*Figure 2-5 Distribution of nearest-station distances among 344 GeoNet strong motion stations*

Figure 2-6 presents an overview of the GeoNet sensor network. The network uses a range of communication links, including cellular, satellite, data radio, broadband and REANNZ (Research and Education Advanced Network New Zealand) (Petersen et al., 2011; REANNZ, 2021). The data is processed, stored and archived in GeoNet servers and cloud servers. GeoNet has two data centres in Auckland and Wellington (Avalon). Amazon Web Services (AWS) are used to deliver earthquake notifications (AWS, 2021). In addition, AWS is used to store data in the cloud. Most urban strong motion sensor stations transmit data using a single cellular link of one of two providers (Spark and Vodafone). Some of the sensors are connected via Spark and others via Vodafone. There is no proper redundancy in transmission links as shown in Figure 2-6. However, if there is an outage in one provider, it will lose connection from those sensors only. Satellite links are primarily connected to an earth station in Sydney, Australia and then data is transmitted to New Zealand data centres. This provides robust connections for remote stations, yet the latency is high compared to cellular and other wired links.

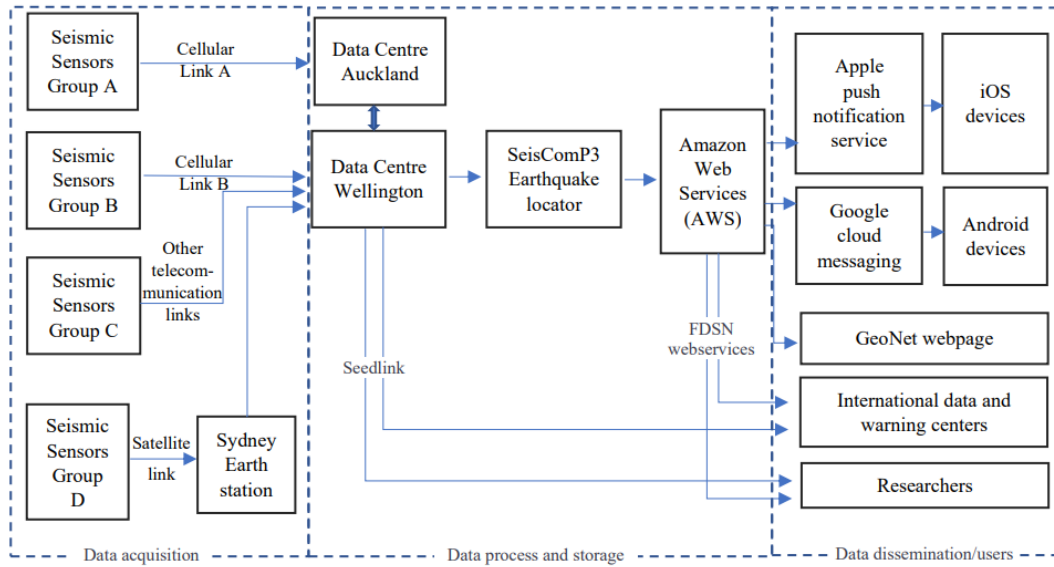


Figure 2-6 Typical GeoNet seismic data flow

The Seedlink protocol is used to transmit data from GeoNet sensors to data centres (Petersen et al., 2011). It is a TCP/IP (Transmission Control Protocol/ Internet Protocol) based real-time data acquisition protocol (GFZ, 2021). GeoNet uses Mini-SEED (Standard for the Exchange of Earthquake Data) format for data transmission (GNS Science, 2024a; Townend et al., 2009). A data packet consists of an 8-byte header and a 512 byte Mini-SEED record (IRIS, 2021). Real-time data could be accessed via the Seedlink protocol and FDSN (The International Federation of Digital Seismograph Networks) web services are used to access near-real-time and archive data (Suárez et al., 2008; GNS Science, 2024a).

## 2.5 EXSIM ground motion simulator

An EEW warning algorithm needs to be tested on a range of earthquake magnitudes, depths, locations, and mechanisms, which will require the shaking waveforms (seismograms) from the earthquakes occurring at multiple locations. However, New Zealand only has limited data on recorded large earthquakes. Therefore, this research generated earthquake seismogram data by simulation. The simulations should address a range of frequencies and require path and site effects.

Several other studies have used synthetic seismograms to evaluate EEW algorithms, including Zollo et al. (2009); Zuccolo et al. (2016); Mittal et al. (2019); McGuire et al. (2021); Thompson

et al. (2024). Zollo et al. (2009) evaluated a probabilistic evolutionary algorithm for EEW for southern Italy, which used p-waves to locate events and estimate magnitudes, and then predicted ground motion intensity at a given target site. They simulated possible rupture scenarios for one M6.9 and two M6 earthquakes. A hybrid k-squared source model is used to compute synthetic seismograms, and a 1-D ground motion prediction equation is used to predict the peak motion at distant sites. Zuccolo et al. (2016) simulated twelve earthquake scenarios using a method (Crempien & Archuleta, 2015) for simulation of broadband ground motion from kinematic earthquake sources to test EEW scenarios at critical facilities in the Eastern Caribbean. Mittal et al. (2019) used the EXSIM program (Motazedian & Atkinson, 2005) to simulate two historical earthquakes in Northern India (Chamoli earthquake of 1999-Mw 6.5 and Uttarkashi earthquake of 1991-Mw 6.8). They used the generated data to test an algorithm called Pd (Wu & Kanamori, 2005), which estimates earthquake magnitude or shaking intensity from the peak displacement of the initial P-wave. McGuire et al. (2021) and Thompson et al. (2024) tested the ShakeAlert EEW system for M9 megathrust earthquakes in the Pacific Northwest using synthetic seismograms. They used a dataset generated from a broadband 3D simulation combining 3D finite-difference simulations and stochastic synthetics (Frankel et al., 2018).

There are several types of earthquake ground motion simulators, and they differ in their complexity, accuracy and computational cost.

- Ground motion prediction equations: These are empirical models based on observed data from past earthquakes to develop shaking metrics as a function of distance and magnitude. These equations are commonly used, in OpenQuake (Pagani et al., 2014; Worden et al., 2017; Böse et al., 2022) and ShakeMap (Cultrera et al., 2013; Worden & Wald, 2016).
- Stochastic propagation models: These models provide reasonable approximations of ground motion characteristics (S-wave arrival time, amplitude and duration, and time history frequency content) for a low computational demand. e.g., SMSIM (Boore, 2005), EXSIM (Atkinson & Assatourians, 2015).
- Physics-based propagation models: The most precise but the highest computational cost. They model the physical processes of earthquake rupture and wave

propagation. e.g., AXITRA (Coutant, 1989), SORD (Ely et al., 2009), Hercules (Tu et al., 2006; Taborda et al, 2010), SPECFEM3D (Komatitsch, & Tromp, 1999).

Therefore, stochastic propagation models are a robust compromise between providing realistic simulation metrics and reducing computational time.

We used EXSIM - a stochastic finite-fault ground motion simulation program (Atkinson & Assatourians, 2015) – to generate synthetic seismograms. EXSIM is not as computationally expensive as some other methods requiring simulations in 2D and 3D media as it is a stochastic-based method. It is a suitable method to use in ground motion based EEW algorithms as it generates realistic time histories in terms of shaking arrival times, durations, and peak amplitudes.

In EXSIM, a fault is modelled as a collection of sub-faults. Sub-faults are considered as point sources. A point-source stochastic model is used to get time series from the collection of sources considering source duration and path duration. Seismic moment and stress parameters are considered to determine the Fourier spectrum at the source. EXSIM also uses a dynamic corner frequency approach to capture realistic frequency ranges from small to large earthquakes. Synthetic seismograms at the observation site are obtained by summing the generated ground motions of each sub-fault.

EXSIM simulates only the S-wave component of the ground motion and calculates the amplitude of the shaking as a single component—it does not generate the three components that are typically recorded by seismographs.

EXSIM requires many parameters to be specified. Some of these parameters need to be calibrated for the particular context, such as according to the region's shear wave (S wave) velocity and seismic wave attenuation model (i.e., Q factor). Chapter 3 describes an analysis in which synthetic seismograms are generated from a recent earthquake and compared to GeoNet recorded data to calibrate the parameters and validate the EXSIM methodology in the New Zealand context.

## **2.6 RNC2 synthetic earthquake catalogue.**

A physics-based realistic synthetic earthquake catalogue for New Zealand was introduced by Shaw et al. (2022), which is especially useful for research work, including this thesis, since NZ has limited recorded data. It is a synthetic catalogue of earthquakes over a 220000-year period to help identify the expected shaking in the future without relying on the short historical record. It is being developed under the same RNC2 earthquake and tsunami project as this thesis.

The active faults (about 530) of New Zealand were reported in the National seismic hazard model of New Zealand (Stirling et al., 2012), and those fault models were used to generate the catalogue. The Rate-and-State Earthquake Simulator (RSQSim), which is based on the approximations of the rate-and-state friction equations (Richards-Dinger & Dieterich, 2012), was used to develop the catalogue by simulating how the faults may develop and cause earthquakes.

Sample catalogue earthquakes from a 500-year period are shown in Figure 2-7. Shaw et al. (2022) stated that this sub-catalogue is spatially consistent with the GeoNet historical catalogue of New Zealand, which began in the 1800s and includes records of a few large events prior to that time. Therefore, the New Zealand RNC2 RSQSim catalogue is a good dataset to test the EEW algorithm by filling the gap of limited recorded data.




Figure 2-7 Sample catalogue earthquakes from a 500-year period. Size of circles indicates relative rupture area (Taken from Shaw et al. (2022), Figure 2-b).

The full catalogue contains 1,609,645 earthquake events, all in the 3.30 to 9.06 magnitude range. The following details of each event are published as a supplement to Shaw et al. (2022):

- Occurrence time.
- Magnitude, moment, and area.
- Number of participating elements, average slip, and long-term average slip rate.
- Hypocentre and centroid locations.
- Rupture surface minimum and maximum depths.

There are a few limitations to this first version of the catalogue. One consequence is that the catalogue more accurately describes the current state and near future than a full prediction for the distant future. However, this limitation is not a problem for testing EEW algorithms since evaluating an EEW system only needs to consider the near future. The other limitation – not including unknown and blind faults – is more significant; for example, the 2011 Christchurch earthquake was caused by a previously unknown blind fault. Using the catalogue for evaluating an EEW system is likely to overestimate the performance of the system because it would ignore the effects of potentially damaging unknown faults, such as the Christchurch earthquake fault. In addition to that, offshore faults are not included in the catalogue except for the Hikurangi subduction zone interface along the east coast of the North Island.

In this thesis, multi-fault earthquakes in the catalogue are modelled as single-fault events, with inferred strike and length parameters. The methodology for this simplification is detailed in Section 4.2.4. Given the existing uncertainties within the catalogue itself, simplified source models are sufficient for analysing EEW system performance in the context of future earthquake scenarios. Similar simplifications have been used in EEW studies by other researchers (e.g., Cremen, Galasso, and Zuccolo, 2022).

# Chapter 3 Preparatory Analysis

This chapter presents several preparatory analyses to prepare for the main analysis presented in chapters 4 and 5. The goal of the first experiment was to determine the average latency time in the GeoNet strong motion sensor network so that the correct latency values could be included in the evaluation of the PLUM EEW algorithm. The second analysis applied the basic PLUM algorithm to recorded seismogram data of two recent earthquakes (Lake Grassmere and Kaikōura) to check that it is plausible to apply PLUM in the NZ context. The third analysis validated parameters for the EXSIM simulator by comparing EXSIM-generated synthetic seismograms to GeoNet-recorded seismograms for the 2010 Darfield earthquake. The fourth analysis explored an initial implementation of the PLUM EEW algorithm on the Darfield earthquake and a set of 209 synthetic earthquakes around the Canterbury region. These analyses confirmed the potential applicability of the PLUM algorithm in the NZ context and identified metrics for evaluating the performance of PLUM.

## 3.1 Data latency of GeoNet sensor network.

A critical part of any networked-based EEW system (i.e., other than onsite systems) is communicating information from the sensors to the central point from which alerts will be issued. Delays due to latency in the communication network will reduce the amount of warning time that can be provided with alerts and must be taken into account in the evaluation of an EEW system.

We conducted an experiment to determine the data transmission latency of the GeoNet strong motion sensor network using a machine located at Victoria University of Wellington (VUW). Locations of the strong motion sensors are shown in Figure 3-1. The VUW machine was connected to the Avalon server and requested a real-time feed of data packets from all strong motion sensors (GeoNet makes such real-time feed available to researchers). Data from the sensors in the network is transmitted as mini-SEED packets to the Seedlink server at Avalon (Wellington), either directly to Avalon or via the Auckland or Sydney data centres. The VUW machine recorded each packet's arrival time and compared it with the time stamp of

the waveform start time in the data packet. The difference in the time stamp is the total latency.

The total latency is a result of several steps:

- $t_p$  = mini-SEED packetisation at the sensor station. Almost all of  $t_p$  is the duration of the signal reported in the packet, with a very small additional time for constructing the data packet.
- $t_n$  = transmission time of the network (cellular, satellite or any other).
- $t_i$  = transmission time on the internet from the network company to the Avalon server, possibly via the Auckland or Sydney data centres.
- $t_d$  = processing of the data at the Avalon server.
- $t_c$  = connection and transmission time between the Avalon server and the VUW computer.

The total latency  $t_{tot}$  is the sum of the components:

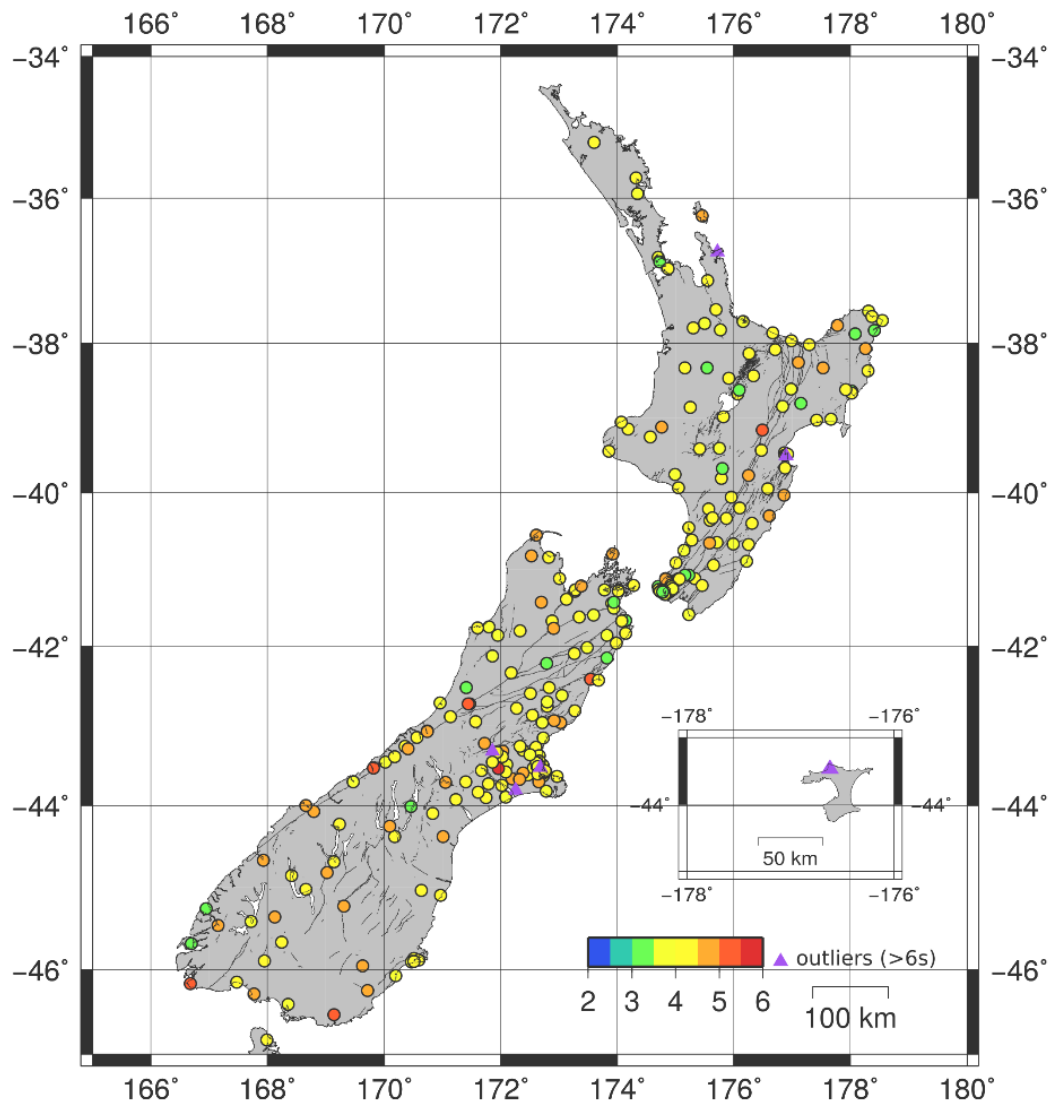
$$t_{tot} = t_p + t_n + t_i + t_d + t_c \quad (3.1)$$

Figure 3-2 shows the distribution of median latencies (without outlier cases) for packets from each individual strong motion sensor using data from one week (19th May 2021 to 25th May 2021). Six sensors had a latency of over 6 seconds, which were considered outliers. They are shown in purple in Figure 3-1 and listed in Table 3-1.

91% of the packets had a total latency between 3.3 and 4.8 seconds. 5% of packets had a latency below 3.3 seconds. A few packets (4%) had a latency above 4.8 seconds. 98% of the data packets were received within 5 seconds. Therefore, we can assume that the current setup can provide data from 98% of strong motion sensors within 5 seconds. This suggests that any EEW system based on the current GeoNet strong motion sensor network setup would have to handle around 5 seconds latency.

It is difficult to precisely decluster the sources of the total latency ( $t_{tot}$ ) from this experiment, since we cannot access the different data transmission stages (e.g., packet arrival times in the communication company networks). However, ping tests between VUW and Avalon suggested that the final component of the latency ( $t_c$ ) is very small (less than 10ms). Individual station data (as shown in Figure 3-1) indicates that higher latencies for remote rural stations

are expected as these are connected via satellite links suggesting that  $t_n$  is a significant component of the latency, at least for remote stations.



*Figure 3-1 Strong motion sensor network - median latency while accessing raw sensor data from VUW via real-time seedlink. Colour of each circle shows the median latency for each strong motion sensor. Outlier stations shown as purple triangles (see Table 3-1).*




Figure 3-2 Distribution of the median latency of the strong motion sensors

Table 3-1 Data latency - outlier cases (latency > 6s)

Station	Median latency (mm:ss)
CTZ	00:07
SBRC	00:34
FAIS	02:20
NCHS	02:47
PRPC	03:49
KOWC	19:22

However, the GeoNet sensor network communication protocols are designed for monitoring and recording, not for EEW. The current packetising mechanism waits to build a data packet containing 400 – 500 samples and may take up to 2 – 2.5 sec, which improves efficiency for monitoring and recording, but is not appropriate for an EEW system. Reducing the data packets to 50 – 100 samples would reduce the packetising time ( $t_p$ ) to well under a second (though it would result in many more packets). Alternatively, if the stations immediately sent very small packets, specifically for EEW, whenever the PGA first exceeded critical thresholds, then the packetising latency would be even lower.

According to Figure 3-2 the current median latency is 4.1 seconds. If the data packets are reduced to 50 samples, it would need to wait approximately 0.25 seconds to collect data at a 200 Hz sampling rate. Changing 450 sample data packets to 50 sample data packets could reduce 2 seconds from the total latency ( i.e.,  $400 \text{ samples} / 200 \text{ Hz} = 2 \text{ s}$ ) bringing the latency down to around 2 seconds ( $4.1 - 2 = 2.1 \text{ s}$ ).

In addition to that, introducing local processing at the sensor data acquisition stage would further reduce packetization time and network traffic, thereby reducing latency even more. Combined with continuous advancements in communication technology, it is reasonable to assume that a future EEW system using the GeoNet strong motion sensor network could achieve an overall latency of just 2 seconds except for some very distant stations that depend on satellite connections.


### **3.2 Initial exploration of PLUM on two recent NZ earthquakes.**

An initial analysis was conducted using GeoNet's recorded earthquake data for two recent major earthquakes—Kaikōura (M7.8, 2016) and Lake Grassmere (M6.6, 2013)—to understand the PLUM algorithm. A simple version of PLUM was used in this analysis, and a few selected GeoNet sensor locations were used as both alerting and receiving points.

In the simple version of PLUM (see Figure 3-3), when a station detects shaking above the defined threshold, it issues an alert. Then, everyone within the warning radius of the station will receive the alert. To determine whether the alert was timely for a station that received the alert, we need to analyse the shaking data at the receiving station to check whether the alert arrived at the receiving station before shaking went above the threshold, and to calculate the warning lead time (Note: communication and processing delays were assumed to be zero, which gives overly optimistic warning times for this initial exploration).

$$\textit{Warning lead time} = \textit{detection time at receiving station} - \textit{detection time at alerting station}$$

To detect shaking above the threshold, the algorithm requires six consecutive samples of the shaking waveform (i.e., 0.03 seconds) to be above the threshold. This was similar to the detection method in the version of PLUM used in Japan (Yamamoto et al., 2008).



*Figure 3-3 Basic PLUM configuration: Alerting station detects shaking above threshold, and issues warning to all points within the warning radius (green circle). The alert can be validated as correct and timely if shaking at some receiving station within the warning radius first detects shaking above threshold after the alert was received.*

GeoNet pre-processed (using SeisComP) strong motion data was used for this analysis. Data processing includes removing instrument gain and high- and low-pass filtering based on the event magnitude. They use signal-to-noise criteria to filter out station data, and therefore, stations with better data for the relevant earthquake were selected for this analysis.

A key parameter for the PLUM algorithm is the peak ground acceleration threshold at which it triggers a warning. In this testing, the threshold was set at  $90 \text{ cm/s}^2$  ( $9.2\%g$ ), which corresponds to the minimum acceleration of MMI VI (Dowrick, 1996; Wald et al., 1999) and is labeled “strong shaking” in the NZ MMI scale (GNS Science, 2023a). This threshold was used as the magnitudes of both of the selected earthquakes are relatively high.

The original PLUM algorithm uses the JMA intensity scale of Japan, which is based on the vector amplitude of three components x, y and z (Hoshiya et al., 2010). This analysis focuses on the S-wave, for which the horizontal acceleration components are higher than the vertical acceleration components. Therefore, to reduce the computational cost of the PLUM algorithm, this analysis only uses the two horizontal acceleration components.

The original PLUM algorithm only considered the sensors within a 30 km radius as mentioned in Section 2.1.1. However, according to the current separation between sensors in the New Zealand GeoNet sensor network, it is difficult to keep within the 30 km radius. Minson et al. (2020) have used sensors within 50 km to test the PLUM EEW algorithm for California. Cochran et al. (2019) have reported good performances for the PLUM algorithm with the

sensors within 60 km for Californian earthquakes. Therefore, the analysis selected validating stations reasonably close to the earthquake epicentre that were not more than 60 km from each other.

For the Kaikōura earthquake (GeoNet ID 2016p858000), three stations were selected as shown in Figure 3-4. Details of the stations are given in Table 3-2. WTMC (Waiau) is where the maximum horizontal PGA was recorded and is the closest station to the epicentre. Both WAKC (in Waikari) and KIKS (in Kaikōura) are within 60 km of the WTMC station.

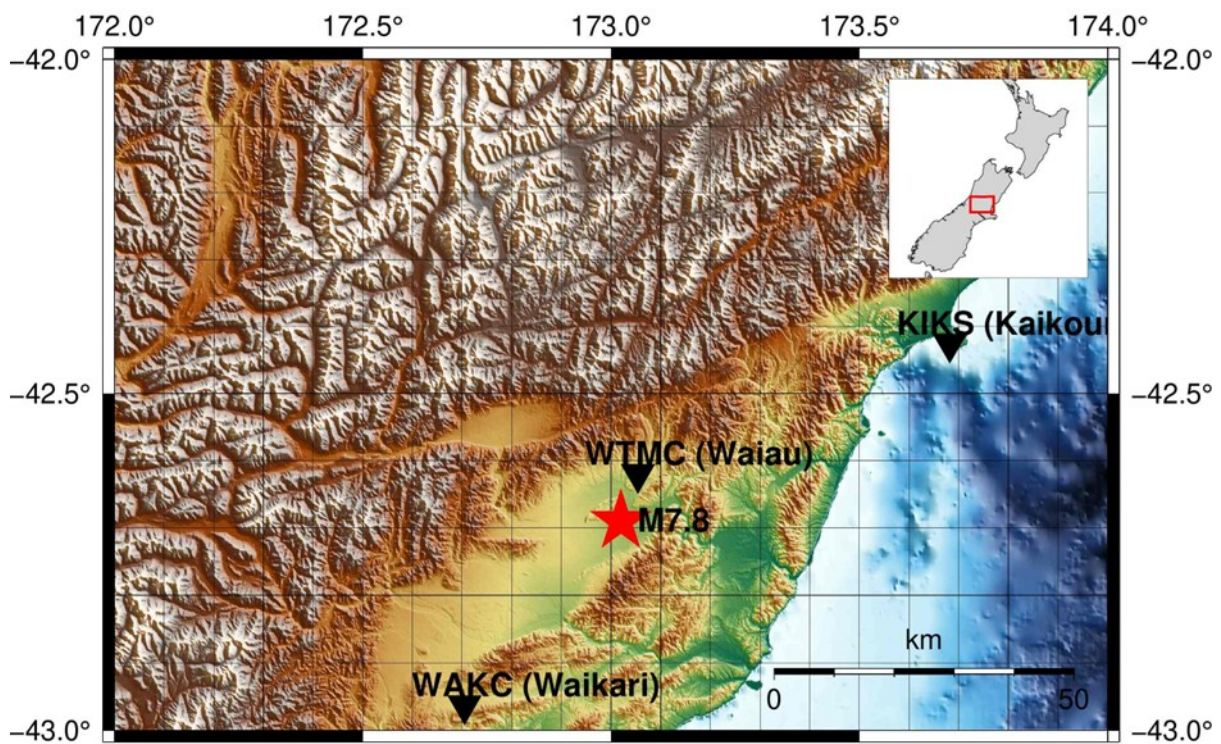


Figure 3-4 Three stations for PLUM simulation with the 2016 Kaikōura earthquake. WTMC provides an alert to WAKC and KIKS.

The PLUM algorithm was applied to the shaking data recorded for the three stations, and the results are shown in Table 3-2. The algorithm detected shaking above the threshold ( $\geq 90$  cm/s<sup>2</sup>; MMI VI) 6 seconds after the origin time at the closest station Waiau (WTMC). According to the analysis, this provides 13 seconds prior warning for the Waikari area and 27 seconds prior warning for the Kaikōura area, where WAKC and KIKS stations are located, respectively.

*Table 3-2 Three stations selected for the Kaikōura earthquake, giving location, distance from earthquake epicentre, distance from previous station, and maximum shaking at the station. PLUM simulation results are the time after the earthquake origin when strong shaking ( $\geq 90 \text{ cm/s}^2$ ) was first detected, and the maximum warning time that the station could have received from the previous station.*

<b>Station</b>	<b>Lat / Lon</b>	<b>Distance from epicentre</b>	<b>Distance from WTMC</b>	<b>PGA max horizontal (<math>\text{cm/s}^2</math>)</b>	<b>Detection Time</b>	<b>Warning time</b>
<b>WTMC</b> (Waiau)	-42.620 / 173.054	9 km	-	1134	6 s	-
<b>WAKC</b> (Waikari)	-42.963 / 172.705	40 km	48 km	147	19 s	13 s
<b>KIKS</b> (Kaikōura)	-42.426 / 173.682	62 km	62 km	256	33 s	27 s

The second earthquake is the 2013 M6.6 Lake Grassmere earthquake (GeoNet ID 2013p613797) using three stations shown in Figure 3-5. Details of the stations are listed in Table 3-3. This earthquake damaged buildings on both sides of Cook Strait, including Seddon and Wellington, which are located 50 km away from the epicentre. As it is a lower magnitude earthquake than the Kaikōura, this exercise was also performed to identify how PLUM will work for that kind of moderate-magnitude earthquake and to understand how an earthquake in the northern part of the South Island could issue EEW alerts to Wellington.

Again, the PLUM algorithm was applied to the shaking data recorded for the three stations, and the results are shown in Table 3-3. Strong shaking ( $\text{PGA} \geq 90 \text{ cm/s}^2$ ) arrived 7 seconds after the origin time at the closest station, Ward, (WDFS) and after 28 seconds, at the Karori station (WNKS) as shown in Table 3-3, and at the Kaitoke station (KIRS) after a further 16 seconds. This indicates that the algorithm would have provided at most 21 seconds warning prior to incoming strong shaking for Karori, Wellington, and at most 16 seconds warning for Kaitoke.




Figure 3-5 Three stations for PLUM simulation with the 2013 Lake Grassmere earthquake. WDFS provides an alert to WNKS, and WNKS provides an alert to KIRS.

Table 3-3 Three stations selected for the Lake Grassmere earthquake, giving location, distance from earthquake epicentre, distance from previous station, and maximum shaking at the station. PLUM simulation results are the time after the earthquake origin when strong shaking ( $\geq 90 \text{ cm/s}^2$ ) was first detected, and the maximum warning time that the station could have received from the previous station.

Station	Lat / Lon	Distance from epicentre	Distance from previous station	PGA max horizontal ( $\text{cm/s}^2$ )	Detection Time	Warning time
<b>WDFS</b> (Ward)	-41.827 / 174.138	14 km	-	591	7 s	-
<b>WNKS</b> (Karori)	-41.285 / 174.742	67 km	79 km	207	28 s	21 s
<b>KIRS</b> (Kaitoke )	-41.077 / 175.230	113 km	47 km	119	44 s	16 s

These warning time results do not take into account communication and processing delays, which would reduce these warning times. Nevertheless, it is possible to use the results from Section 3.1 to consider the data latency. The measurements from Section 3.1 suggest that the total data latency  $t_{\text{tot}}$  would be about 5 seconds, using the current networking parameters.

Assuming the processing delay is negligible, then the maximum warning times shown in Table 3-2 and Table 3-3 will be reduced by up to 6 seconds. The warning times might be reduced by only 2 seconds if the changes to the networking parameters suggested in Section 3.1 were implemented. Note, these times do not include the delivery time to end users.

Only a few stations are used for this preliminary testing, and most are around 60 km apart. The analysis showed, using recorded data of real earthquakes, that implementing the PLUM algorithm for New Zealand is plausible, but the analysis is far too limited to give a reliable indication of its performance on the expected range of earthquakes that might occur in the near future.

### **3.3 Evaluating EXSIM on NZ earthquake data.**

We could improve the testing of the PLUM algorithm on recorded earthquake data by expanding the analyses discussed in Section 3.2 to more earthquakes. However, testing the PLUM algorithm for a range of earthquakes for an appropriate range of stations and grid locations is not possible using the limited recorded data. Recorded data is available for the sensor locations only, whereas simulations could generate data for a range of locations and various scenarios. The EXSIM simulator, described in Section 2.5 was selected to generate synthetic seismograms.

First, an actual earthquake was selected to simulate using the EXSIM and the synthetic waveforms were compared against the recorded data to validate the EXSIM simulation. This analysis was useful for calibrating the appropriate parameters for EXSIM and confirming that EXSIM is an appropriate tool to generate synthetic seismograms for a wide range of earthquakes in NZ for the purposes of testing an EEW algorithm for New Zealand.

#### **3.3.1 Synthetic seismograms of Darfield earthquake**

The 2010 M7.1 Darfield earthquake was selected to generate synthetic seismograms using the EXSIM simulator and to compare them with the real recorded data. The 2010-2011 Canterbury earthquake sequence, which began with the Darfield earthquake, was the most impactful seismic event series in New Zealand since World War II (Berryman, 2012). Given its significance, the Darfield earthquake was selected as a useful starting point for analysis.

The Darfield earthquake is particularly interesting and useful because it was a relatively complex event involving multiple faults (Gledhill et al., 2010), and a rich recorded dataset is available (Cousins & McVerry, 2010). This earthquake and its aftershock sequence (Townend et al., 2012) provided New Zealand's most extensive set of strong motion data since recording began in the early 1960s (Cousins & McVerry, 2010). It also served as the most significant test of GeoNet since its establishment in 2001 (Gledhill et al., 2010). Additionally, fault models for this earthquake have been developed by researchers (Holden et al., 2011; Beavan et al., 2012), which are valuable for simulating the earthquake.

To run the EXSIM simulation, appropriate parameters had to be selected based on the fault models of the Darfield earthquake. Key parameters included the S-wave velocity, the fault configuration, and the appropriate hypocentres.

To choose the S-wave velocity, we considered the model of Ristau (2008) with suggested S-wave velocities for various levels in both the North Island and the South Island. Following that model, an S-wave velocity of 3.7 km/s used in the simulations. The frequency range for the simulations was set between 0.05 Hz and 10 Hz.

Initially, we used the fault model from Holden et al. (2011), simulated using EXSIM. However, that model did not include the initial fault near the epicentre. Therefore, the model was further refined by including data from Beavan et al (2012), and then adjusting parameters to generate seismograms reasonably compatible with the recorded GeoNet seismograms. This exercise was useful for adjusting the EXSIM parameters.

### **Details of Fault model**

We simulated the two main fault segments named Charing-cross and Greendale with the parameters given below. The two main fault segments come from the seismic model of Holden et al. (2011). Also, we have included a fault near the epicentre (the fault which initially ruptured) using the model of Beavan et al. (2012). EXSIM requires slip data for the simulation. Modelled slip data used in Holden et al. (2011) is used for the Charing-cross and Greendale segments as reported in GNS Science (2017). Beavan et al. (2012) uses slip data based on GPS data which includes slip from aftershocks and therefore overestimates the slip during the fault rupture and is therefore unreliable for the ground motion simulation. Therefore, the

simulation used random slip distribution for the initial fault. The fault segments are shown in Figure 3-6 and the parameters shown in Table 3-4 were used for the simulation.

The simulation results which were generated using the fault model, agreed with the observed data reasonably well (and better than the results of an earlier model we tested. The results are described in the following section.

*Table 3-4 Parameters of fault model*

<b>Fault</b>	<b>Rupture start time (s)</b>	<b>Magnitude</b>	<b>Strike</b>	<b>Dip</b>	<b>Fault type</b>	<b>Fault size (length x width)</b>	<b>Hypo location (along-fault, down-dip distance)</b>
A. Initial fault near epicenter	0.00	6.2	150°	54°	strike-slip	9 km x 9 km	4 km, 4 km
B. Charing Cross	3.00	5.8	040°	65°	reverse	20 km x 20 km	10 km, 3 km
C. Greendale	6.63	6.9	266°	82°	strike-slip	40 km x 20 km	30 km, 10 km

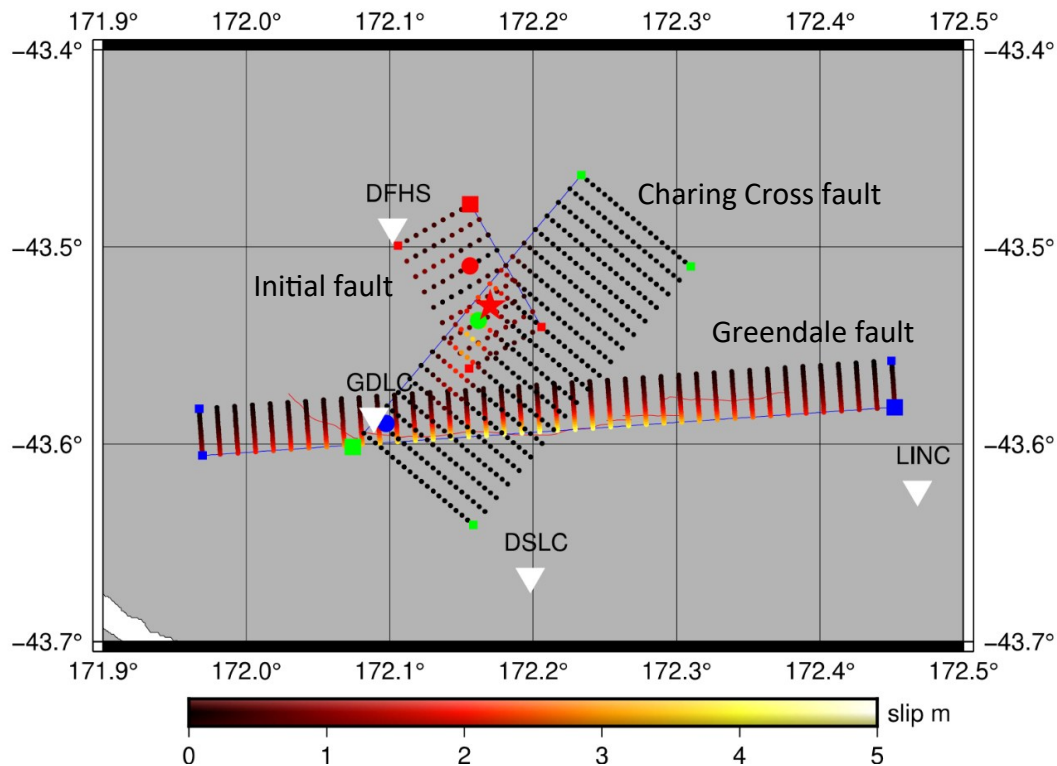


Figure 3-6 Darfield earthquake fault segments and slip distribution – fault model.

The circles represent the hypocentre of each sub-fault and the squares represent the four corners of the faults (red: fault (a), green: fault (b), blue: fault (c)). The larger squares are the EXSIM ref point (upper corner) of each fault. The blue lines are the tops of the faults. The rows of small coloured circles projecting from the tops of the faults show the slip distribution at each point down the fault according to the colour scale where black=0m, yellow=5m. The red star is the GeoNet hypocentre of the event. The white inverted triangles are the stations.

### 3.3.2 Comparison of the synthetic and GeoNet data

The data recorded by the 32 GeoNet stations within 127 km of the epicentre was compared with the synthetic data generated by the simulation. The 127 km limit was chosen because these 32 stations include the GeoNet stations that experienced  $14 \text{ cm/s}^2$  or greater shaking.  $14 \text{ cm/s}^2$  (i.e., 1.4% g) is considered as the threshold for the MMI IV shaking (Wald et al., 1999). According to the GeoNet strong motion data (GNS Science, 2022) these 32 stations have the best signal-to-noise ratio.

We first present a summary of the key differences between the synthetic data and GeoNet data across all the stations, then give more detailed comparisons of the synthetic and GeoNet seismograms for three of the stations.

Figure 3-7 **Error! Reference source not found.** shows the locations of the stations, and uses colour to show the PGA differences between Synthetic and GeoNet data, relatively as percentage with more intense colours indicating greater differences.




Figure 3-7 Comparison of Synthetic and GeoNet PGA data. Circles represent the station locations, circle colour represents the difference between the PGA in synthetic and GeoNet data ( $(Syn-GeoNet)/GeoNet$ ), blue: synthetic PGA < GeoNet PGA; red: synthetic PGA > GeoNet PGA.

The colours of the circles show the relative differences in PGA. There are a few stations around the faults where the synthetic PGA values are higher than the GeoNet data, but most of the circles are light blue or light red and are considered acceptably close results. There are a few underestimated synthetic PGA results (dark blue) in the east towards Christchurch City, probably caused by the direction of the earthquake rupture. The PGA values are shown as a

scatter plot in Figure 3-8 **Error! Reference source not found.**, and it is indicated that the three stations close to the epicentre have a slightly higher synthetic PGA compared to the GeoNet values. However, for the other stations, the synthetic PGA is on the lower side and within a reasonable range.

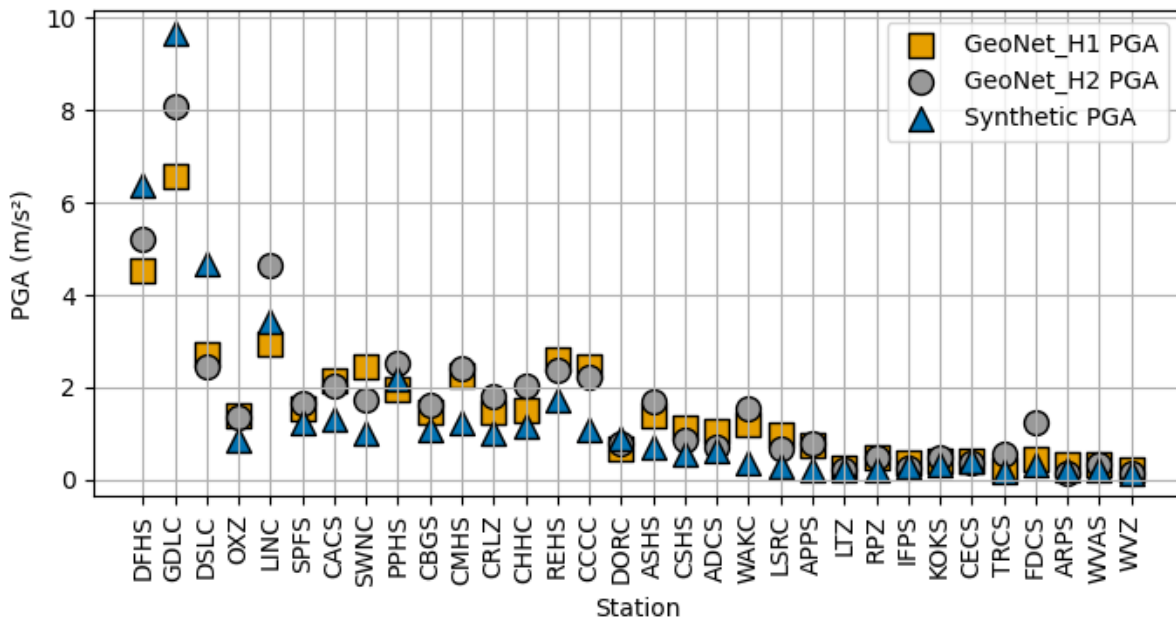


Figure 3-8 Comparison of modeled (synthetic) and observed (GeoNet) PGA Values. The data recorded by the 32 GeoNet stations within 127 km of the Darfield earthquake epicentre was compared with the synthetic data generated by the simulation

The purpose of the simulation data is to evaluate the PLUM algorithm. A critical value in PLUM is the level of shaking that triggers the detection, so we need to determine whether both synthetic and actual data triggered the detection at the same time or not. At this point, we are using the arbitrary detection threshold of  $50 \text{ cm/s}^2$ . This threshold is within category V MMI, and is a reasonable initial value, but later analyses will explore other thresholds.

The colour of the triangles in Figure 3-9 **Error! Reference source not found.** represent the detection time difference between the synthetic and GeoNet data for the PLUM algorithm, calculated with the  $50 \text{ cm/s}^2$  threshold. For most of the stations that detected shaking, the synthetic detection time is slightly after, or the same as, the GeoNet detection time (less than 3 seconds). For one station (SPE), the synthetic detection time is more than 5 seconds after the Geonet detection time, and for two stations to the northeast of the hypocentre, the synthetic detection time is a few seconds before the GeoNet detection time.

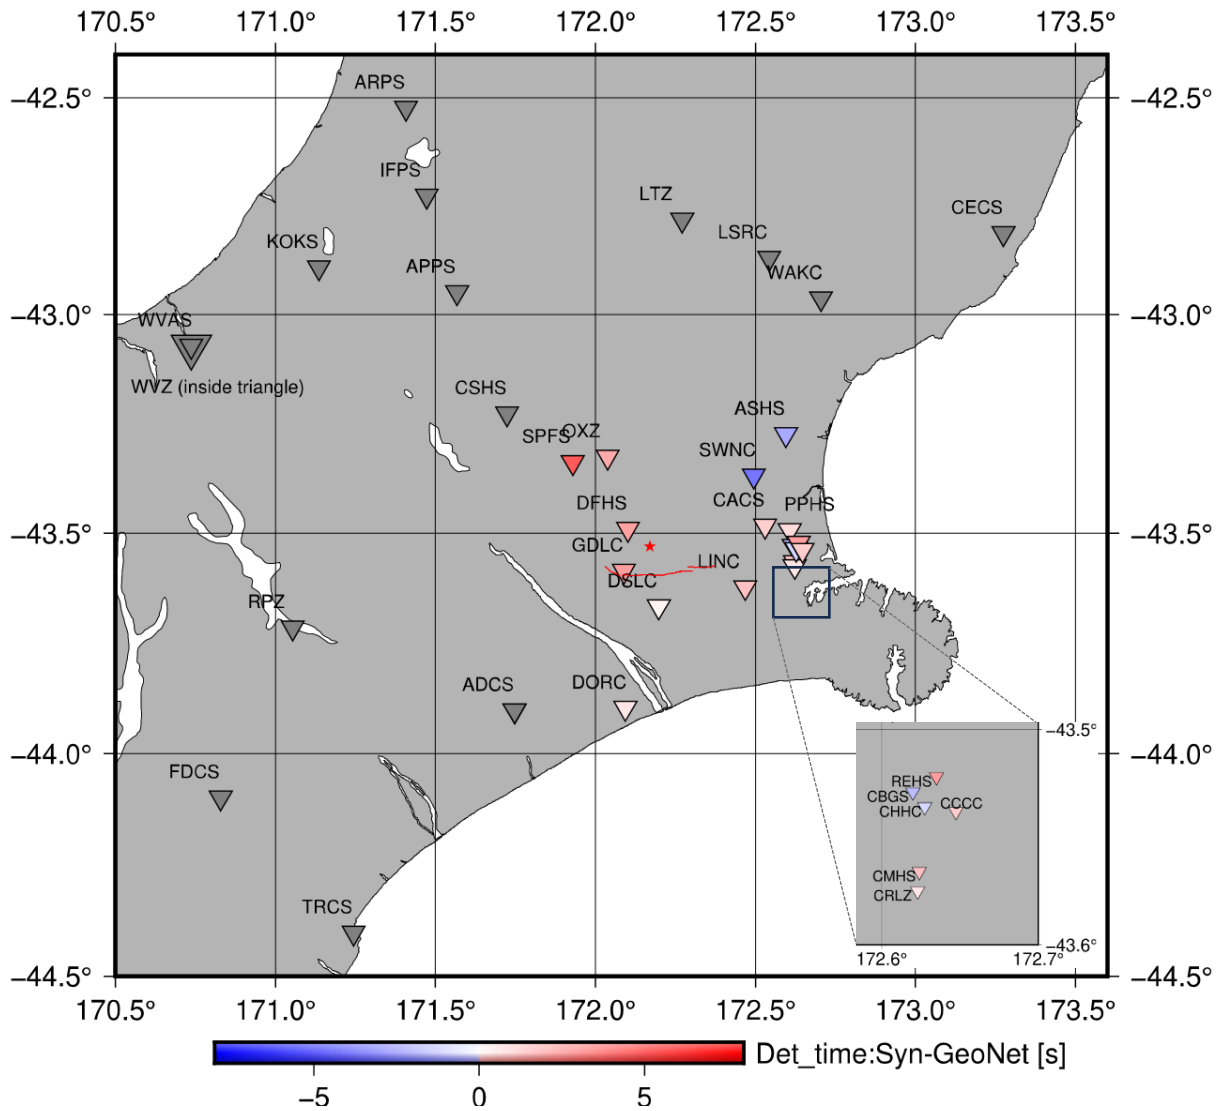


Figure 3-9 Comparison of Synthetic and GeoNet detection time data. triangles represent the station locations, triangle colour represents the time difference between detection of 50 cm/s<sup>2</sup> shaking in synthetic and GeoNet data (blue: synthetic before GeoNet; red: synthetic after GeoNet). Grey triangles indicate stations where the shaking did not reach 50 cm/s<sup>2</sup>.

*Figure 3-10* shows the synthetic seismograms and actual GeoNet data for the DFHS station which is located on the initial fault segment and also within 7 km of the epicentre. Synthetic seismograms of each of the three separate fault segments and the combination of the three seismograms are shown in the figure along with two horizontal channels of GeoNet recorded data. The first line illustrates the synthetic seismogram for the initial fault segment near the epicentre. PGA and detection time of the shaking (the time to exceed the PLUM threshold) of each waveform are also mentioned in the label above the waveform. The second line is the synthetic seismogram of the Charing-cross fault, and the third is the Greendale fault. The fourth line shows the combined synthetic seismogram for all faults for this earthquake and was used for the testing of the PLUM algorithm. The last two lines illustrate the GeoNet recorded waveforms from the two horizontal channels H1 and H2. The maximum PGA and the minimum detection time of the two channels were used for the comparisons with the synthetic data.

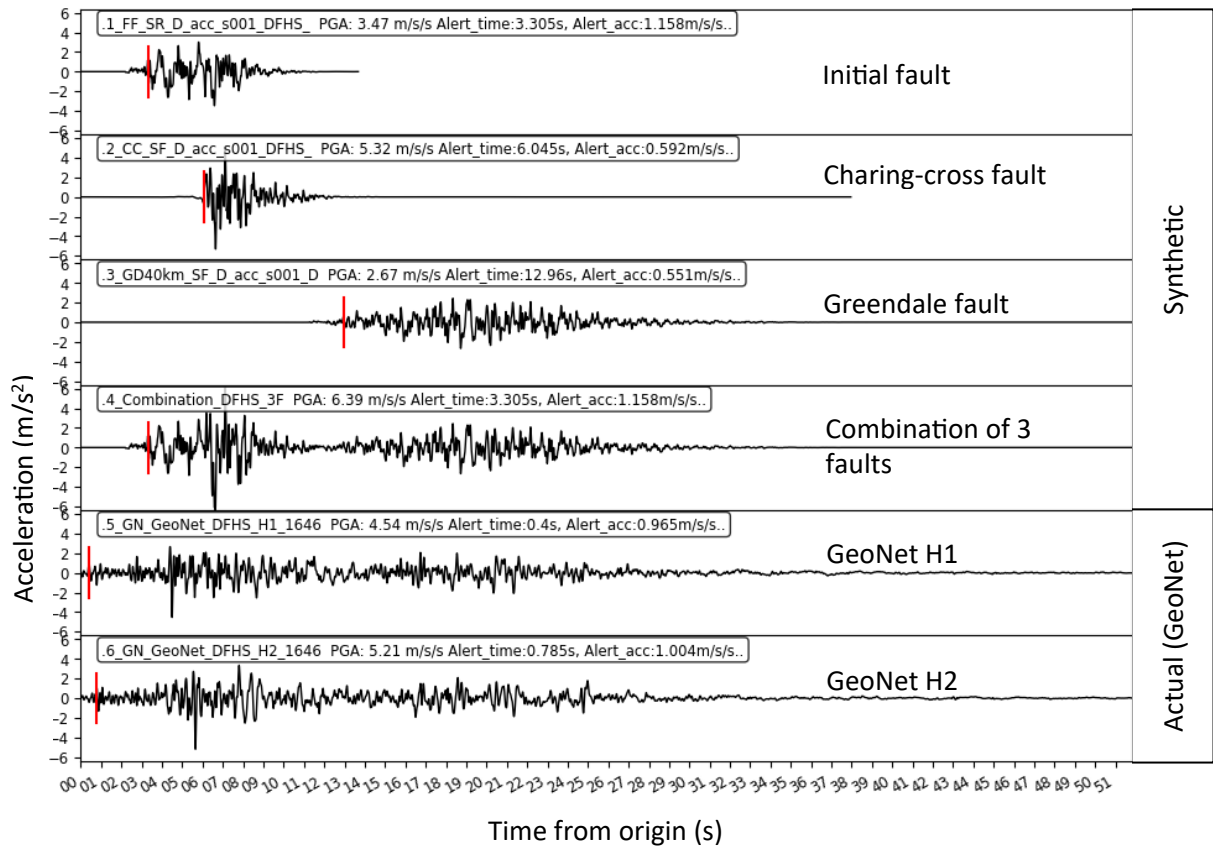


Figure 3-10 Comparison of the Synthetic and GeoNet seismograms: DFHS station. Detection times of the shaking (the time to exceed the PLUM threshold) are indicated by red vertical lines. Note that the Synthetic seismograms are single channel, showing the amplitude of the S-wave shaking; the GeoNet seismograms show that the two horizontal components of the shaking.

In the DFHS station, the PLUM alert is activated by GeoNet data within one second of the earthquake origin time. For the synthetic data, the alert is generated at 3.3 seconds. The reason for this difference in detection time may be the unavailability of P-wave data in the synthetic seismograms. As mentioned in Section 2.5, the EXSIM ground motion simulation produces S-wave data only and does not generate P-wave data. The data from the GeoNet stations, which were located around the fault/epicentre may have sufficient P-wave component on the horizontal component to hit the threshold. Therefore, some of the real stations near the fault triggered the alert for GeoNet data earlier than for the synthetic data. Figure 3-11 shows the synthetic seismograms and actual GeoNet data for the CACS station, which is further from the faults. The distance between the CACS station and the epicentre is 30 km. Detection times of the shaking (the time to exceed the PLUM threshold) for both synthetic and actual data are consistent, as shown in the figure. The alert detection time for the synthetic seismogram is 13.14 seconds. Recorded seismograms of GeoNet

sensor's two horizontal channels H1 and H2 triggered the alert at 11.64s and 14.12s, respectively. The latter part of the GeoNet waveforms (after 30 seconds) is not included in the synthetic seismogram, as we didn't simulate the faults which were ruptured after the Greendale fault. We did not include those latter faults as our focus is on the initial major faults which will trigger the PLUM algorithm.




Figure 3-11 Comparison of the Synthetic and GeoNet seismograms: CACS station. Detection times of the shaking (the time to exceed the PLUM threshold) are indicated by red vertical lines. Note that the Synthetic seismograms are single channel, showing the amplitude of the S-wave shaking; the GeoNet seismograms show that the two horizontal components of the shaking.

Figure 3-12 shows the synthetic and actual waveforms of the DORC station, which is located 41 km away from the epicentre. The synthetic waveform and two GeoNet horizontal waveforms are reasonably equal within the first part of the signals. According to Figure 3-11 and Figure 3-12, the generated synthetic seismograms are adequately matched with the actual GeoNet waveform data and can be used for the PLUM algorithm testing.

Simulating the ground motion from an actual earthquake to generate synthetic seismograms similar to the recorded data is challenging as it depends on the various parameters including

the fault model and the simulator. Nevertheless, these results validate the EXSIM simulation as an appropriate tool to test the PLUM algorithm for synthetic earthquakes.

The Darfield earthquake was simulated considering the published fault model and parameters. As it is a multi-fault rupture, it was complex to simulate using EXSIM. To evaluate PLUM for a wide range of earthquakes, it is unfeasible and computationally expensive to simulate multi-fault ruptures. Therefore, in the later analyses, earthquakes are assumed to be single-fault ruptures and simulated accordingly.




Figure 3-12 Comparison of the Synthetic and GeoNet seismograms: DORC station. Detection times of the shaking (the time to exceed the PLUM threshold) are indicated by red vertical lines. Note that the Synthetic seismograms are single channel, showing the amplitude of the S-wave shaking; the GeoNet seismograms show that the two horizontal components of the shaking.

### Synthetic earthquake propagation / PGA

We compared the GeoNet and synthetic seismograms for GeoNet sensor station locations, as discussed in the previous section. In addition to the sensor station locations, we also generated synthetic seismograms for a region of Canterbury by dividing the region into

0.1°×0.1° (≈11 km×8 km) grid cells. Synthetic seismograms were generated using EXSIM for the Darfield earthquake source parameters (fault model as described in Section 3.3.1), considering the middle of each grid cell as the target point. The simulated PGA values for each grid point are shown in Figure 3-13. PGA values within 0 to 2 m/s<sup>2</sup> are illustrated by a yellow to brown colour scale, and all PGA values which are greater than 2 m/s<sup>2</sup> are illustrated by a black colour. The maximum PGA of 6.43 m/s<sup>2</sup> was computed for a grid point (lat: -43.55, long: 172.25), which is located on the Greendale fault. In addition to that, Figure 3-14 illustrates the GeoNet shaking layers map (MMI intensity contours) for the Darfield Earthquake.

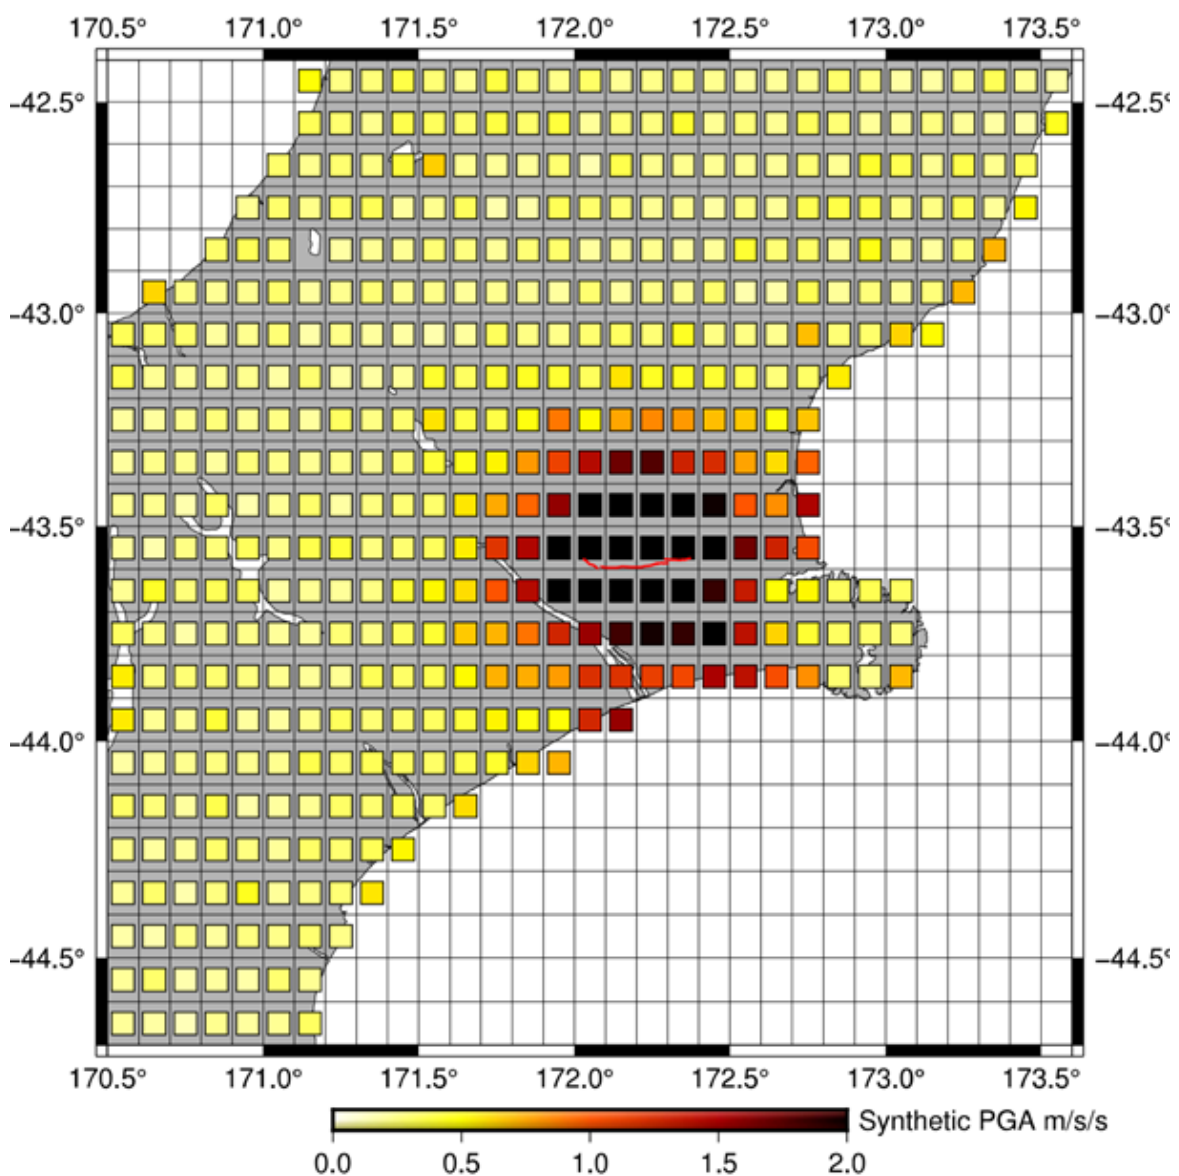


Figure 3-13 PGA of Synthetic seismograms at grid points across Canterbury. Colour of squares shows PGA.  $PGA \geq 2$  shown in black. The Greendale Fault is marked by a red line.

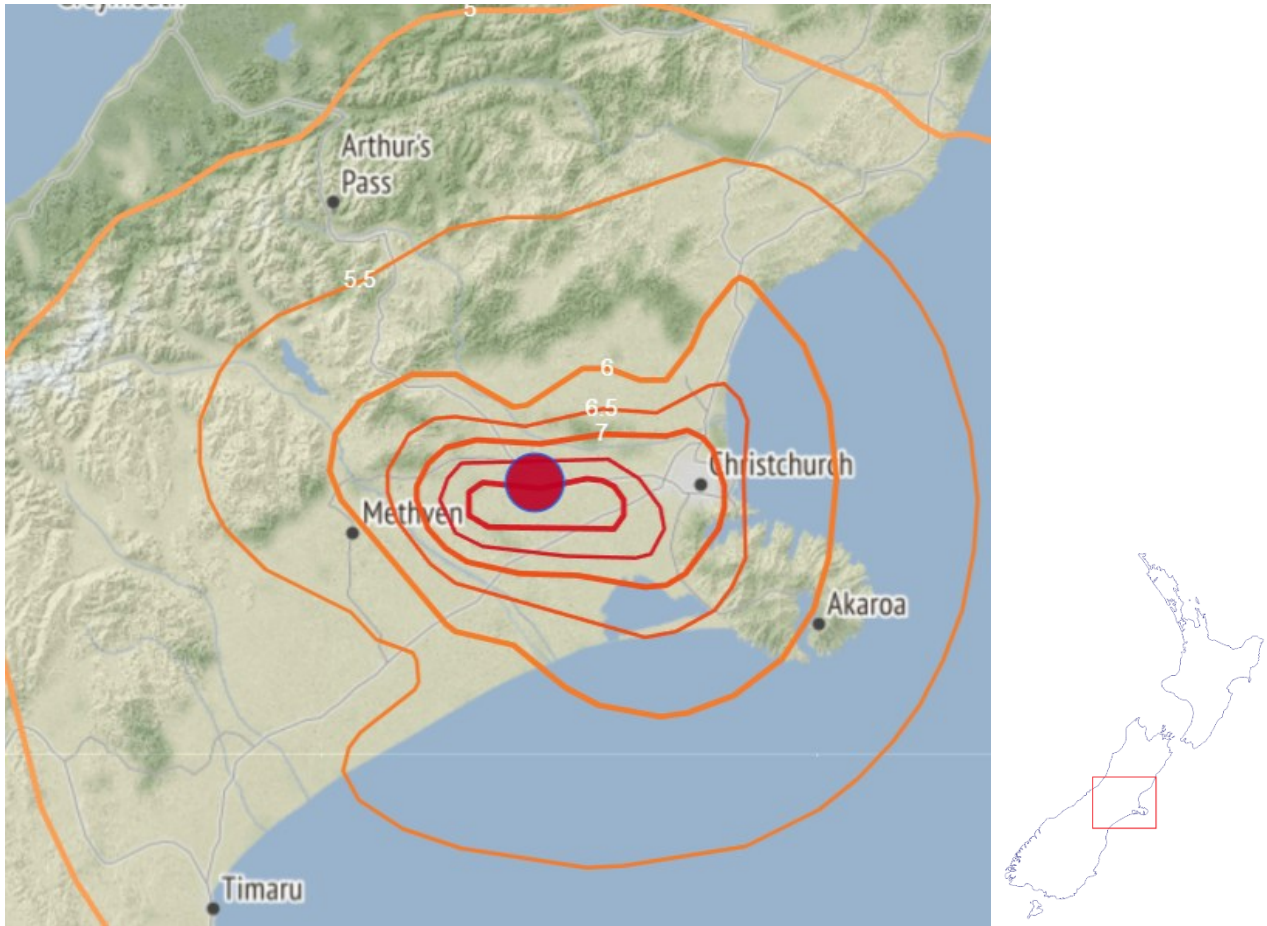


Figure 3-14 GeoNet Shaking Layers Map for the Darfield Earthquake, showing MMI intensity contours. Adapted from "Darfield earthquake," by GeoNet, (n.d.), GeoNet (<https://www.geonet.org.nz/earthquake/felt/3366146>)

### 3.4 Testing PLUM with 209 synthetic earthquakes in the Canterbury region

To explore the use of the PLUM algorithm on synthetic data, we performed a test of the basic version of PLUM using 209 virtual earthquakes, placed in 209 different locations across Canterbury. Canterbury was chosen as the initial region of study because it has a richer collection of actual datasets available than other regions, especially from the Canterbury earthquake sequence.

As in Section 3.3, a region of Canterbury within  $171.1^{\circ}$  to  $173.0^{\circ}$  longitude and  $-43.0^{\circ}$  to  $-44.1^{\circ}$  latitude was divided into  $1^{\circ} \times 1^{\circ}$  ( $\approx 11 \text{ km} \times 8 \text{ km}$ ) grid cells, giving 209 grid locations on land for the epicentres of the synthetic earthquakes. An additional border of  $0.6^{\circ}$  (6 grid cells), with no earthquakes, gave a larger region of 459 grid locations on land within  $170.5^{\circ}$  to  $173.6^{\circ}$  longitude and  $-42.4^{\circ}$  to  $-44.7^{\circ}$  latitude. Figure 3-15 illustrates the region and the grid cells.

The border area was chosen to include all stations within 60km of the central region, since these stations could provide alerts to grid cells in the central region if the PLUM warning radius is at most 60 km. 33 GeoNet station locations (strong motion sensors) were used within the larger region.




Figure 3-15 Canterbury region with grid containing 459 grid cells. The 209 earthquakes are located in the grid cells containing a red star. The labeled black triangles are the locations of the sensor stations.

The properties of the 209 similar virtual earthquakes (earthquakes are shown in Figure 3-15) were chosen to understand/evaluate the PLUM algorithm for the region, considering the following factors.

- magnitude 6.5: large enough to cause damage, but not so large that all parts of the region would experience MMI V or above (PLUM threshold) shaking, therefore, it is a better test of PLUM.
- Single fault: As described in Section 3.3.2, it is unnecessarily complex to handle multiple faults, and defining a single fault for the earthquake magnitude is a reasonable choice.
- Simple strike-slip shallow fault: typical of earthquakes in that region that can cause damage.

First, the ground motion from all 209 earthquakes was simulated using the EXSIM parameters listed in Table 3-5. For each earthquake, synthetic seismograms for all 459 grid points and for all 33 stations were generated. All the grid points were assumed to have site class D, but the sensor stations were allocated the site classes listed in Kaiser et al. (2017b).

*Table 3-5 EXSIM parameters for simulation.*

<b>Parameter</b>	<b>Value</b>
Magnitude	6.5
Stress	9 MPa
Fault size	10 x 10 km
Fault type	Strike slip
Strike	0°
Dip	90°
Depth of top of the fault	1 km
Hypocentre location along-fault	5 km
Hypocentre location down-dip distance	5 km
Slip pattern	Random
S-wave velocity	3.7 km/s
Synthetic seismograms	S-wave only

The 102,828 synthetic seismogram data sets (459\*209+ 33\*209) were then used to test the PLUM algorithm. For each earthquake, synthetic seismograms at each station were used to determine when the station would first detect strong shaking (above 38 cm/s<sup>2</sup>) and therefore issue an alert and synthetic seismograms at each grid point were used to determine the arrival time of strong shaking at the grid point. The threshold of 38 cm/s<sup>2</sup> was chosen as it is at the

lower level of MMI V, which can cause some level of damage (GNS Science, 2023a). The timeliness of an alert was computed by comparing the alert receiving time and the arrival time, taking the communication delay into account. As in Section 3.2, to determine the arrival time of the strong shaking, the algorithm looked for a sequence of 6 samples of the acceleration (0.03 s for the 200 Hz data, which were above the threshold of  $38 \text{ cm/s}^2$ ). This was modelled on the example of the Japanese system, which used a sliding window on the sequence of samples to determine the shaking level.

### 3.4.1 Performance for different warning radii.

Performance of the PLUM algorithm for different warning radii was evaluated using the generated synthetic seismogram data and the following performance measures for each grid cell.

- CTA (correct timely alerts): number of alerts received before strong shaking arrived.
- LA (late alerts): number of alerts received after strong shaking arrived.
- MA (missed alerts): number of times strong shaking arrived but no alert was received although there was at least one station within the PLUM warning radius.
- MANS (missed alerts due to low station density): number of times strong shaking arrived but no alert was received because there was no station within the PLUM radius.
- FA (false alerts): number of alerts that were received, but no strong shaking arrived.
- NANS (no alerts, no shaking): the number of times no alert was received, and no strong shaking arrived.
- MLA (missed or late alerts): LA+MA+MANS

The communication and processing delay was assumed to be 2 seconds (see Section 3.1) in determining correct timely alerts vs late alerts. Correct timely alerts (CTA) enable precautionary actions (or automated actions) before the shaking arrives; late or missing alerts (LA, MA, MANS) do not allow precautionary actions, so they are grouped together as MLA (missed or late alerts).

Because some grid cells receive more shaking events than others, it is important to use normalized values to compare the performance at different grid cells. Normalized CTA for a

grid point is the fraction of strong shaking events for which a correct timely alert was received.  $CTA_{norm}$  is obtained by dividing the CTA for a grid point by the number of strong shaking events received by the grid point, as shown in equation 3.1. Similarly, the normalized  $MLA_{norm}$  is shown in equation 3.2.

$$CTA_{norm} = \frac{CTA}{No\ of\ shakings \geq Th} = \frac{CTA}{CTA+LA+MA+MANS} \quad (3.1)$$

$$MLA_{norm} = \frac{LA+MA+MANS}{No\ of\ shakings \geq Th} = \frac{LA+MA+MANS}{CTA+LA+MA+MANS} \quad (3.2)$$

For the same reason, the normalized FA for a grid cell is the fraction of non-strong shaking events for which a false alert was received.  $FA_{norm}$  is obtained by dividing FA by the number of earthquakes which did not produce strong shaking (i.e., shaking < threshold) at the given grid location, as shown in equation 3.3. Similarly, the normalized  $NANS_{norm}$  is obtained by equation 3.4.

$$FA_{norm} = \frac{FA}{No\ of\ shakings < Th} = \frac{FA}{FA+NANS} \quad (3.3)$$

$$NANS_{norm} = \frac{NANS}{No\ of\ shakings < Th} = \frac{NANS}{FA+NANS} \quad (3.4)$$

Figure 3-16 and Figure 3-17 show the geographical distribution of  $CTA_{norm}$  and  $FA_{norm}$  performance of the PLUM for all 209 earthquakes with 30 km PLUM radius and 38 cm/s<sup>2</sup> (MMI V) shaking threshold (Wald et al., 1999). The purple rectangles show the inner region with the locations of the 209 earthquakes.

In Figure 3-16 the size of each coloured square represents the number of strong shakings received by that given grid location (the largest square represents 27 shakings, and the smallest square represents 2 shakings). The colour of each square shows the fraction of alerts that were correct and timely ( $CTA_{norm}$ ). Dark blue means 100% correct timely alerts are received by the given grid point. Dark red means 0%  $CTA_{norm}$  (in other words 100%  $MLA_{norm}$ ). A light color between pink and blue represents 50% correct timely alerts.

In Figure 3-17, the size of each coloured square represents the number of times the grid location did not receive strong shakings (the inverse of Figure 3-16), and the colour shows the fraction of those events that received a false alert. For much of the region (dark blue), there

were very few false alerts, but the fraction is higher around Christchurch City, especially immediately to the west of the city.

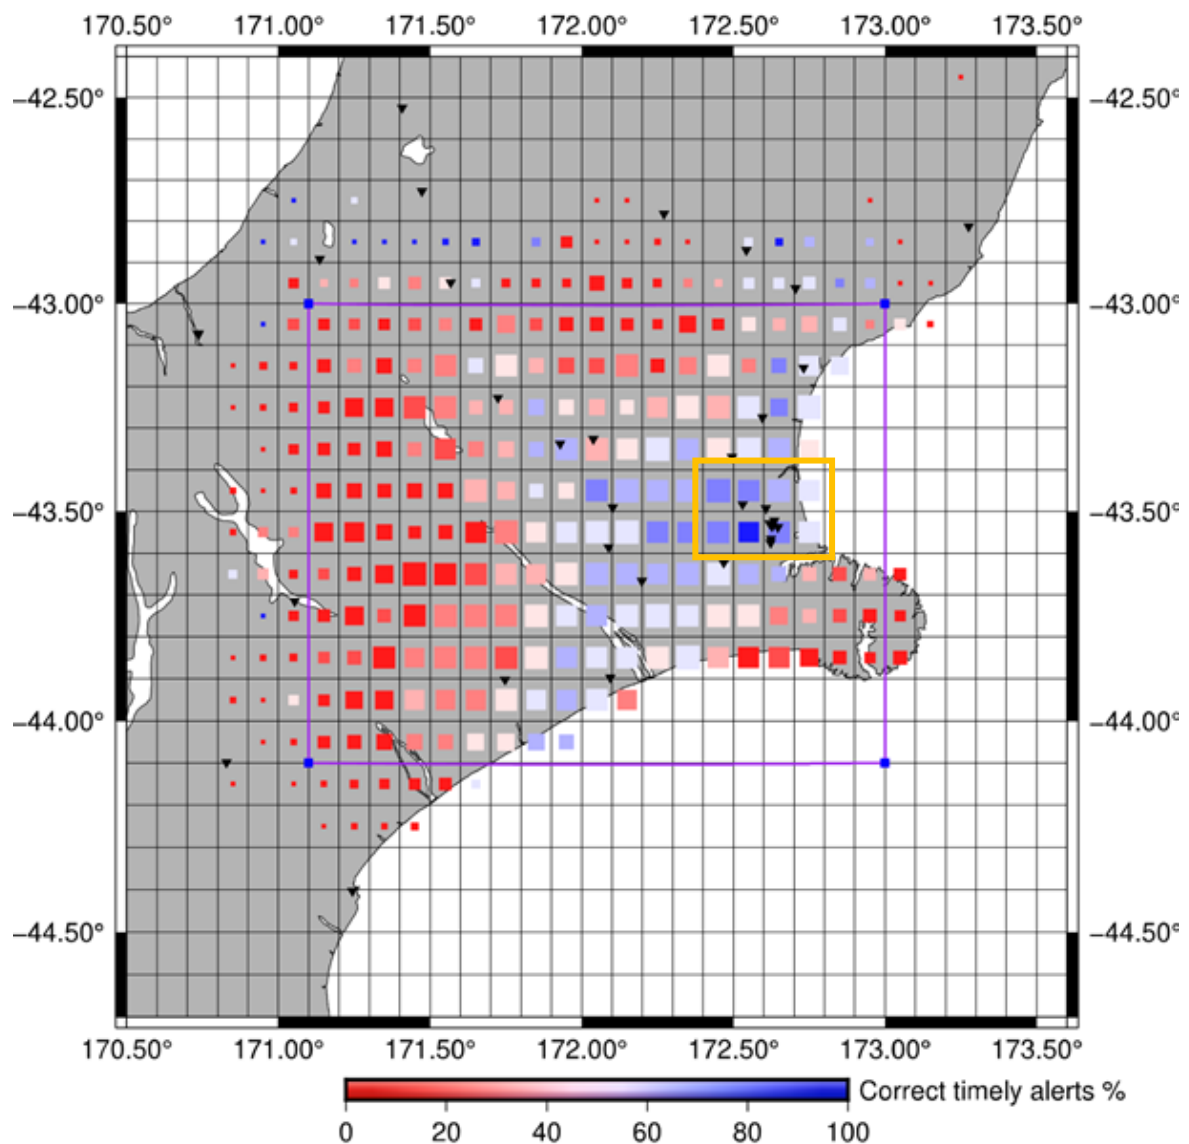


Figure 3-16 Fraction of Correct Timely Alerts for 209 synthetic earthquakes in the Canterbury region. PLUM radius=30 km, threshold=38 cm/s<sup>2</sup>. The purple rectangle shows the inner region with the locations of the 209 earthquakes. The labeled triangles are the sensor stations. The orange rectangle indicates the grid cells covering Christchurch City. Square size represents the number of strong shaking events at the grid cell (the largest square represents 27 shakings, and the smallest square represents 2 shakings). Square colour represents  $CTA_{norm}$ : the fraction of alerts that were correct and timely. Blue = 100% CTA, red = 0% CTA, a light color between pink and blue = 50% CTA.

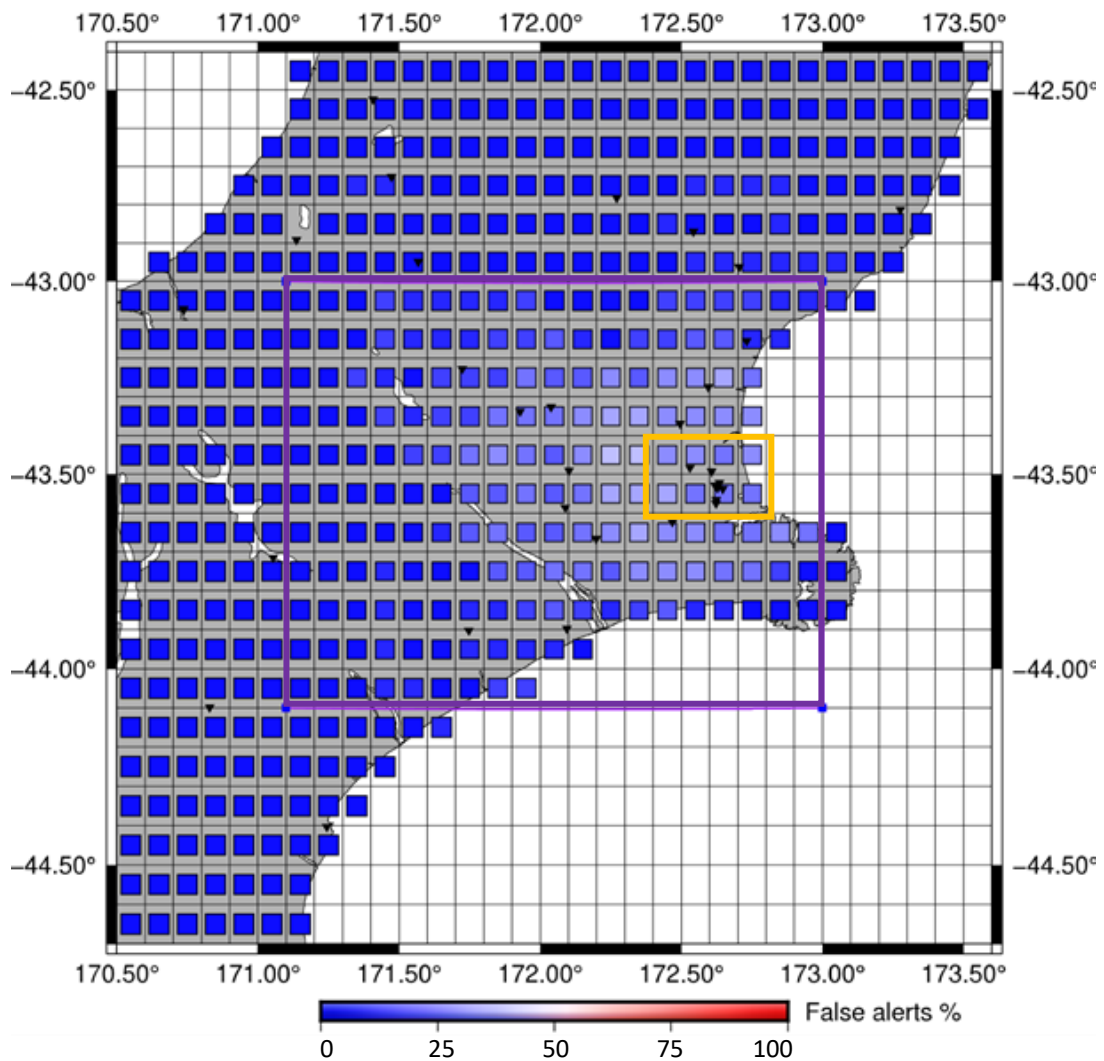


Figure 3-17 Fraction of False Alerts for 209 synthetic earthquakes in the Canterbury region. PLUM radius=30 km, threshold=38 cm/s<sup>2</sup>. The purple rectangle shows the inner region with the locations of the 209 earthquakes. The small white triangles are the sensor stations. The orange rectangle indicates the grid cells covering Christchurch City. Square size represents the number of earthquakes that did not generate strong shaking events at the grid cell. Square colour represents  $FA_{norm}$ : the fraction of those events for which there were false alarms. Blue = 0% FA; red = 100% FA.

To evaluate the performance of PLUM at different warning radii, we summed the measures of the performance across all the grid cells in the region, and compared them at different warning radii. There were a total of 3,430 strong shaking events across all the grid cells, and 92,501 events where the shaking never reached the strong shaking threshold.

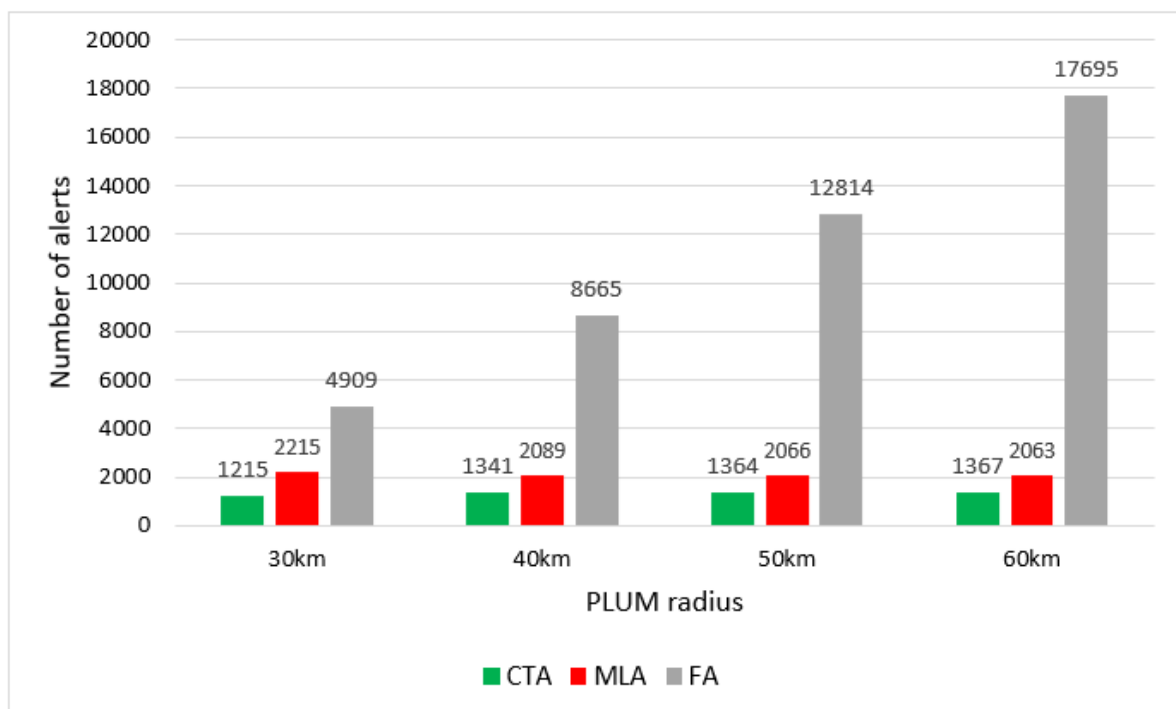
Table 3-6 shows the totals of the basic performance indicators for 30 km, 40 km, 50 km and 60 km radii. The column for CTA/(CTA+MLA) shows the percentage of all strong shaking events for which there was a correct timely alert and the column for FA/(FA+NANS) shows

the percentage of all non-strong shaking events for which there was a false alarm. Figure 3-18 shows a graph of the values of CTA, MLA and FA. Table 3-7 gives the values of the normalised metrics  $CTA_{norm}$  and  $FA_{norm}$  averaged over all the grid cells at the different warning radii.

It is clear that both the number and fraction of correct timely alerts increases with the warning radius, but the number (and fraction) of false alarms also increases more rapidly.

*Table 3-6 PLUM performance over whole Canterbury region, basic metrics. Entries give the total of each metric over all grid cells. 209 earthquakes, threshold=38 cm/s<sup>2</sup>. PLUM radius of 30 km, 40 km, 50 km, 60 km.*

PLUM radius	CTA	LA	MA	MANS	MLA	CTA / (CTA+MLA)	FA	NANS	FA/(FA+NANS)
30 km	1215	1150	736	329	2215	35%	4909	87592	5%
40 km	1341	1276	787	26	2089	39%	8665	83836	9%
50 km	1364	1318	748	0	2066	40%	12814	79687	14%
60 km	1367	1330	733	0	2063	40%	17695	74806	19%



*Figure 3-18 Total of CTA, MLA and FA at different PLUM radii.*

Table 3-7 PLUM performance over whole Canterbury region, normalised metrics. Entries give the average of the normalized metrics over all grid cells. 209 earthquakes, threshold=38 cm/s<sup>2</sup>. PLUM radius of 30 km, 40 km, 50 km, 60 km.

PLUM radius	Average CTA <sub>norm</sub>	Average FA <sub>norm</sub>
30 km	29%	6%
40 km	32%	10%
50 km	33%	14%
60 km	33%	20%

Figure 3-16 and Figure 3-17 suggest that PLUM with a warning radius of 30 km would work well around Christchurch City, providing correct timely alerts for almost all the earthquakes that generated strong shaking near the city and a small number of false alerts. However, PLUM does not appear to work as well away from Christchurch City where there are far fewer sensors, especially in the western half of the region. The summary results illustrated in Table 3-6 and Figure 3-18 show that increasing the warning radius improves the number of correct timely alerts, but at a cost of increasing the number of false alerts. There was a significant 10% improvement in CTA by increasing the radius from 30 km to 40 km, though the false alert rose by 76%. Beyond 40 km, there is no significant improvement on CTA but a continued noticeable rise in false alerts.

The next two sections explore ways of adding new sensor stations to improve the performance of PLUM in this region.

### 3.4.2 Adding new stations.

One of the important parameters of the PLUM algorithm is the warning radius, because a station will only issue an alert to locations within that radius – if the radius is 30 km, a sensor station may issue the alert to the grid points within the 30 km and not beyond that – and therefore far away grid points may have to wait until a sensor station within 30 km detects the earthquake. If there is a grid point that has no stations within the warning radius, then it can never receive a warning.

The analysis in Section 3.4.1 used 33 actual GeoNet sensor locations (the stations with good data for the Darfield earthquake). If there are some grid points that do not have any stations within 30 km, this would contribute to the poor performance of PLUM in some parts of the

Canterbury region. The station density of the selected region was analysed and the results shown in Table 3-8. 99 grid points have no station within 30 km, which reduces to only 4 grid points at 50 km. Within the central region (the purple rectangle containing the locations of the 209 earthquakes), only 22 grid points have no station within 30 km, reducing to 2 grid points at 40 km. Clearly, this will limit the number of correct timely alerts that are possible for a warning radius of 30 km or 40 km. Note that even if every grid point was within the warning radius of at least one sensor, correct timely alerts are still not guaranteed since the relative locations of the earthquake, the station and the grid point will affect whether the station is able to issue a timely alert to the grid point so that adding stations is likely to also improve the performance with warning radii of 50 km and 60 km.

*Table 3-8 Sensor station density in the Canterbury region.*

<b>Distance from grid to closest station</b>	<b>Grid points having no station within radius. (Out of all 459 points)</b>	<b>Grid points in central region having no station within radius. (Out of 182 points in central region)</b>
> 30 km	99	22
> 40 km	30	2
> 50 km	4	0
> 60 km	1	0

We performed an analysis to explore how many new stations would need to be added to ensure that there is at least one station within 30 km of each grid point in the central region, and then re-evaluated the PLUM algorithm with the four different warning radii. The following algorithm was used to find the new station locations to improve the density:

**Algorithm to add stations.**

- While** minimum grid-to-station distance for all grid points in central region < 30 km:*
- Calculate minimum grid-to-station distances for each grid point in central region.*
- Find the grid point with the maximum grid-to-station distance.*
- Place new station at this grid point.*

The algorithm found the five new station locations shown in Figure 3-19. These stations were sufficient to ensure that each grid point within the central region would have at least one station within 30 km.

We then ran the PLUM analysis again, including those five new stations, and the performance of PLUM improved. The differences in the  $CTA_{norm}$  before and after adding new five stations for a warning radius of 30 km are illustrated in Figure 3-20 (a) and (b). The improvement in performance can be seen particularly in the reduced intensity of the red squares near the new stations, especially on the west and north of the region, after adding stations. However, the differences do not stand out clearly.




Figure 3-19 Five new station locations, shown as yellow triangles, that improved station density to ensure that all grid points in the the central region are within 30 km of at least one station. The existing thirty-three GeoNet stations are shown as black triangles.

The improvement can also be seen in the metrics totaled over the whole region as shown in Table 3-9 and Table 3-10, which tabulate the results for warning radii from 30 km to 60 km. At every warning radius, the total CTA increased and the total MLA decreased, but the total FA also increased. The average  $CTA_{norm}$  also increased quite substantially. With the added

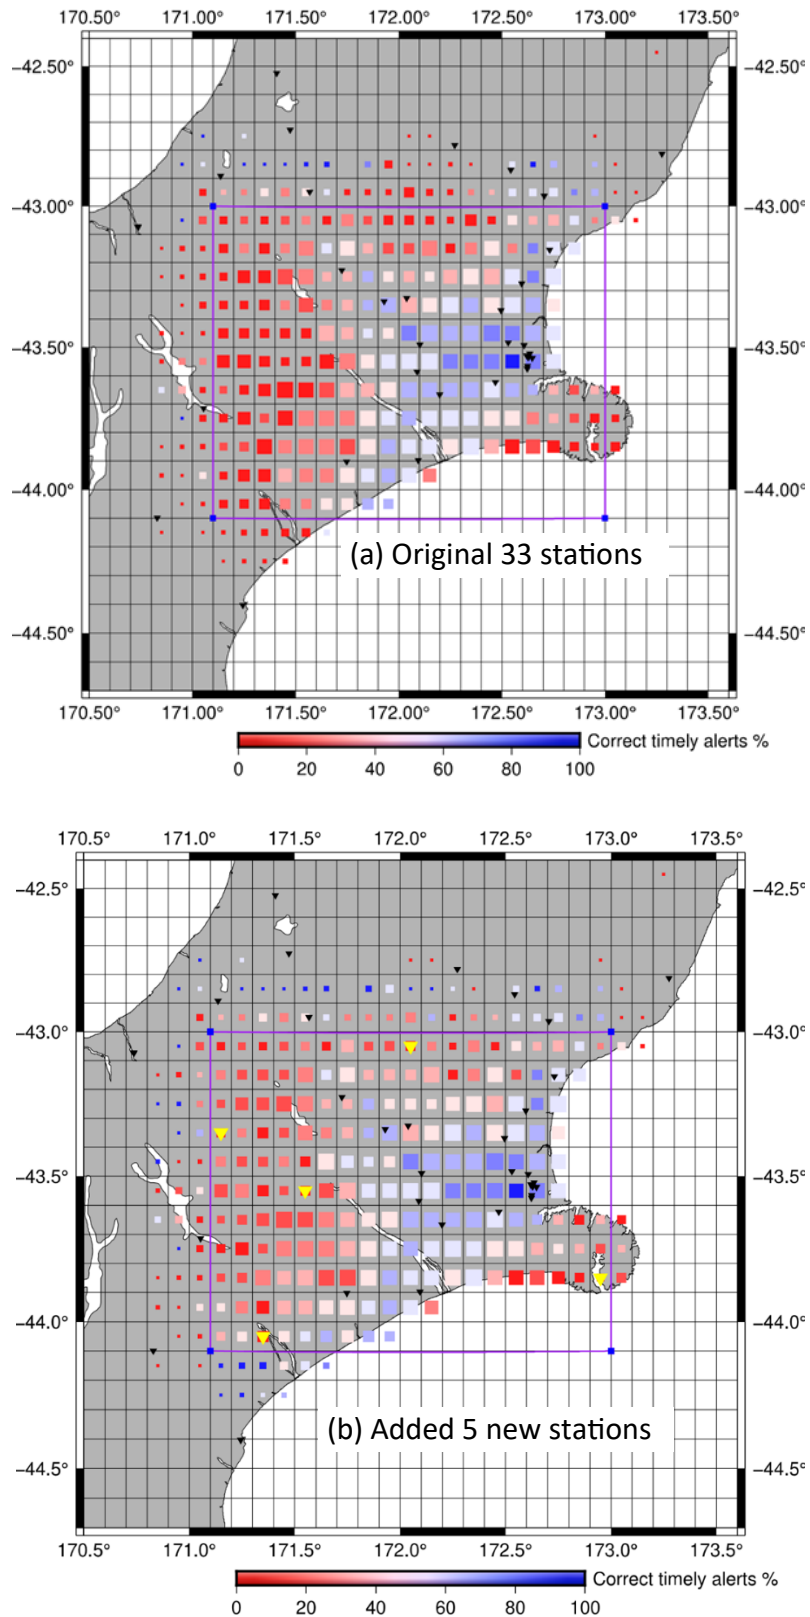


Figure 3-20  $CTA_{norm}$  for warning radius 30 km after adding 5 stations. (a) (same as Figure 3-16), with only the original 33 stations, marked as small black triangles; (b) with 5 additional stations, marked with yellow triangles. As in Figure 3-16, size of squares represents the number of strong shaking events in the grid cell and colour of squares represents the value of  $CTA_{norm}$ .

stations, the total CTA and average  $CTA_{norm}$  were a bit higher at a warning radius of 40 km than 30 km, but increasing from 40 km to 50 km or 60 km made less difference. A more detailed analysis of the changes to the MLA (= LA+MA+MANS) shows that the MANS component reduced to 0 (by design) and the number of missing alerts decreased, but the number of late alerts increased: adding stations increased the number of alerts that could be given, but around half of the new alerts appear to be late alerts or false alerts.

*Table 3-9 PLUM performance before and after adding 5 new stations: total CTA (correct timely alerts), total MLA (missing and late alerts), and average  $CTA_{norm}$  for the original 33 stations, for the 38 stations, and the percentage change.*

Warning Radius	Total CTA			Total MLA			Average $CTA_{norm}$		
	Original	Added	Change	Original	Added	Change	Original	Added	Change
30 km	1215	1408	+16%	2215	2022	-9%	29.0%	39.1%	+35%
40 km	1341	1583	+18%	2089	1847	-12%	32.4%	44.2%	+37%
50 km	1364	1618	+19%	2066	1812	-12%	33.3%	46.1%	+39%
60 km	1367	1622	+19%	2063	1808	-12%	33.3%	46.2%	+39%

*Table 3-10 PLUM performance before and after adding 5 new stations: total FA (false alerts) for the original 33 station, for the 38 stations, and the percentage change.*

Warning Radius	Total FA			Average $FA_{norm}$		
	Original	Added	Change	Original	Added	Change
30 km	4909	5629	+15%	5.6%	6.4%	+14%
40 km	8665	10143	+17%	9.8%	11.4%	+16%
50 km	12814	15254	+19%	14.4%	17.1%	+19%
60 km	17695	21242	+20%	19.8%	23.7%	+20%

It is clear that there was a significant improvement in early warning alerts and that improving station density improves the performance. However, the algorithm uses a greedy approach that finds a local optimum solution for each iteration—the grid point that is furthest from its nearest station—and does not necessarily provide the global optimum solution. It would be computationally expensive to find the global optimum solution and this was not considered in this analysis. Furthermore, the algorithm only considers the distance of grid points from

stations and does not take into account the effect on PLUM performance or that the optimum locations for new stations may depend on other factors such as population density.

The algorithm above improved the station density by adding enough stations to ensure all grid points were within 30 km of at least one station but ignored performance. To address the problem of ignoring the effect on performance of adding stations, we conducted a further analysis to determine the best new station location for obtaining optimal PLUM performance.

In this analysis, a station was added separately at each grid point within the central region and the performance of PLUM (30 km warning radius) was measured in each case with two different measures—total CTA and average  $CTA_{norm}$ . The maximum total CTA was achieved by including a new station at a grid point in the south-west of the region at (lat: -43.85, long: 171.45) as shown by the yellow triangle in Figure 3-21. Adding this station improved the total CTA by 8.7%, as shown in Table 3-11. The maximum average  $CTA_{norm}$  was achieved by including a new station at a grid point further to the south-west, at (lat:-44.05, long: 171.15), as shown by the green triangle in Figure 3-21. Adding this station improved the average  $CTA_{norm}$  by 12.7%.

The analysis was repeated at warning radii of 40 km, 50 km and 60 km. When optimising for total CTA, the optimal new station locations for these warning radii was the same as for the 30 km radius, but the gain in CTA was lower than for 30 km radius. When optimising for average  $CTA_{norm}$ , the optimal new station location for the higher radii was different from the 30 km radius, and the gain in  $CTA_{norm}$  was higher. The changes in performance when optimising for total CTA or average  $CTA_{norm}$  at each warning radius are shown in Table 3-11.

*Table 3-11 PLUM performance when adding one new station at a location optimized for total CTA or a location optimized for average  $CTA_{norm}$ . Shows the total CTA and average  $CTA_{norm}$  before and after adding the station and the percentage change.*

Warning radius	Total CTA (Optimising for Total CTA)			$CTA_{norm}$ (Optimising for Average $CTA_{norm}$ )		
	Original	Added	Change	Original	Added	Change
30 km	1215	1321	8.7%	29%	33%	12.7%
40 km	1341	1453	8.4%	32%	39%	17.5%
50 km	1364	1478	8.4%	33%	40%	16.9%
60 km	1367	1481	8.3%	33%	40%	16.9%

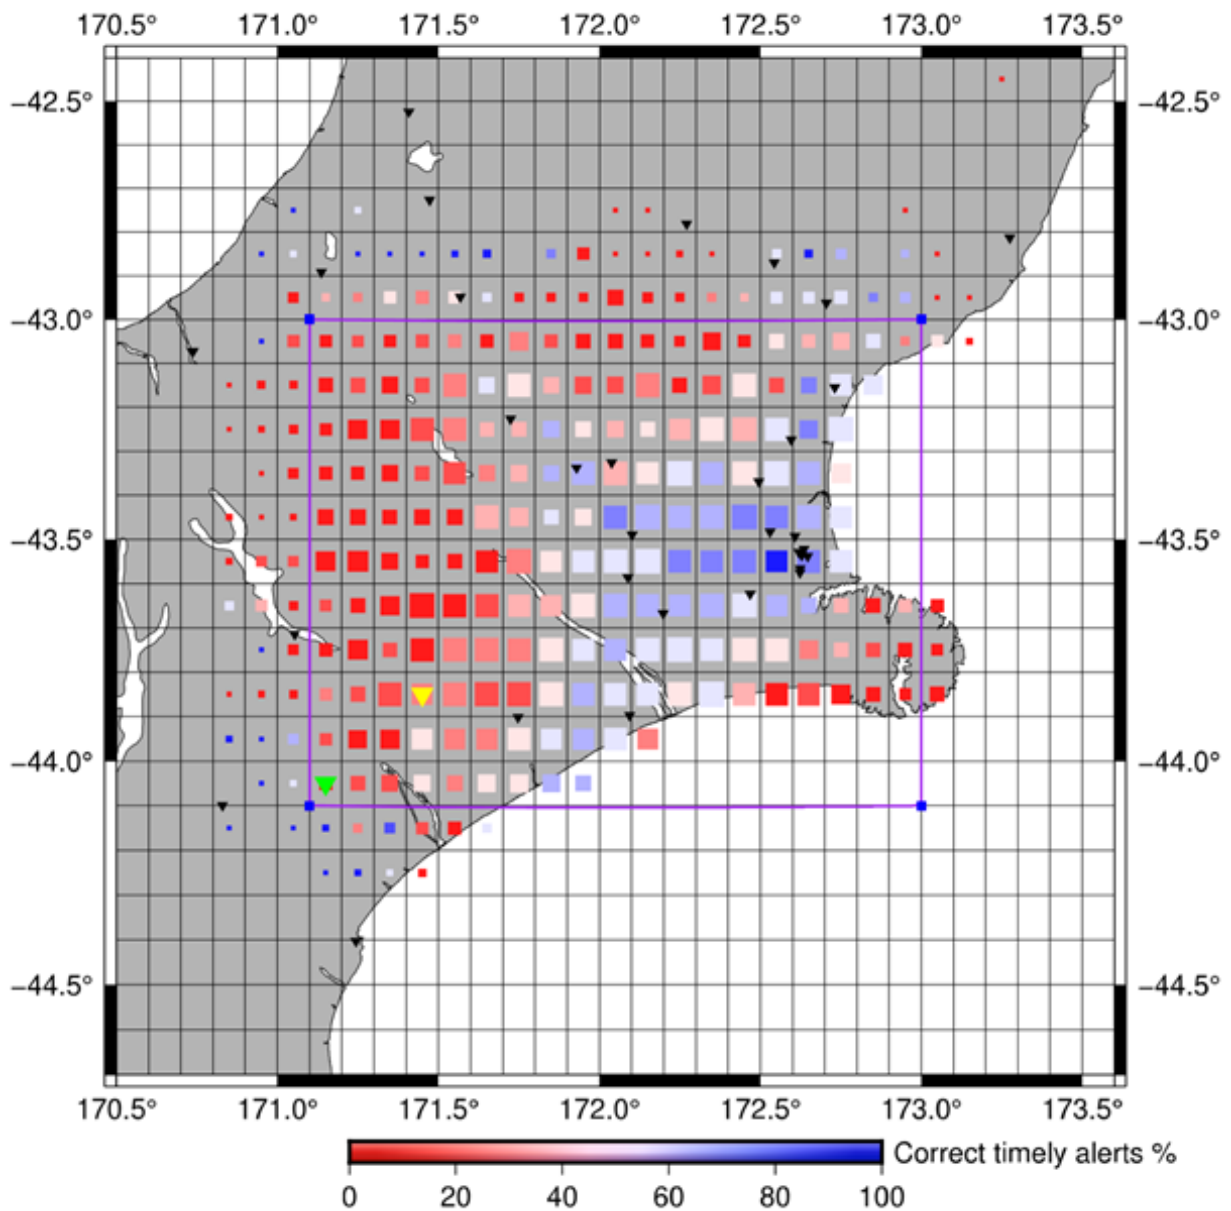


Figure 3-21 Optimal position for a single new station for a warning radius of 30 km considering the  $CTA_{norm}$ . Green triangle shows location for maximising the improvement in the average  $CTA_{norm}$  across the region. As in earlier figures, the colour of the squares shows the value of  $CTA_{norm}$  at each grid point. (Yellow triangle shows location for maximising the improvement in the total CTA across the region).

Comparing the results in Table 3-9 and Table 3-11, with a 30 km warning radius, adding one station at the optimal location for total CTA gives over a half of the gain of adding five stations that minimise distance from grid points to stations, and adding one station at the optimal location for average  $CTA_{norm}$  gives over one third of the gain of adding five stations that minimise distance from grid points to stations.

These results strongly suggest that when choosing a new location for a station, it is useful to consider the effect on the performance of the EEW system rather than just seeking to improve the station density.

This section is only focused on adding stations to improve the performance of a PLUM-based EEW system, showing how the results of the simulation of the PLUM system could be used to contribute to the placement of new stations. The section does not address the broader issue of determining the best locations for new stations in general, for which there is an extensive literature (Böse et al., 2022b; Kuyuk & Allen, 2013; Oth et al., 2010; Petersen et al., 2011; Rawlinson et al., 2012) and many more considerations than are addressed here.

### **3.4.3 Population density and PLUM for Canterbury region**

A problem with the analyses in the previous sections is that they ignored the population distribution. The main goal of an EEW system is to protect people and infrastructure (Becker et al., 2020a). Therefore, the distribution of the population is an important factor for analysing the outcomes of earthquake early warning algorithms. This section takes the population density of New Zealand into consideration to analyse the results of the PLUM algorithm.

We obtained population data from WorldPop (2020). The spatial resolution of this data is 30 arc-seconds ( $0.0083^\circ$ ), giving an area of approximately  $0.62 \text{ km}^2$  for each data cell. We used the data to calculate an average population density in each grid cell.

Figure 3-22 shows the population density across the Canterbury region. The colour of each square shows the average population density of the grid cell on a logarithmic scale where yellow is the highest and black is the lowest. Gray grid cells mean the population density is less than 1 per WorldPop data cell. Note that for the cells on coastline, the population density is averaged just over the land component of the cell. The population is mainly concentrated in Christchurch, with two smaller concentrations to the south (Ashburton and Timaru), but a very low population in the western part of the region.

It is clear that the performance of PLUM in the high density grid cells, mainly Christchurch, is much more significant than the performance in the low density grid cells, especially in the west. To make the performance measures reflect this, the key performance measures—Average  $\text{CTA}_{\text{norm}}$  and Average  $\text{FA}_{\text{norm}}$ —should be weighted averages of the  $\text{CTA}_{\text{norm}}$  and

$FA_{norm}$  values of each grid cell, where the values are weighted by the population density of each grid cell, as in equations 3.5 and 3.6

$$\text{Population-weighted } CTA_{norm} = \frac{\sum_{g=1}^{459} (CTA_{norm_g} \times \text{population density}_g)}{\sum_{g=1}^{459} \text{population density}_g} \quad (3.5)$$

$$\text{Population-weighted } FA_{norm} = \frac{\sum_{g=1}^{459} (FA_{norm_g} \times \text{population density}_g)}{\sum_{g=1}^{459} \text{population density}_g} \quad (3.6)$$

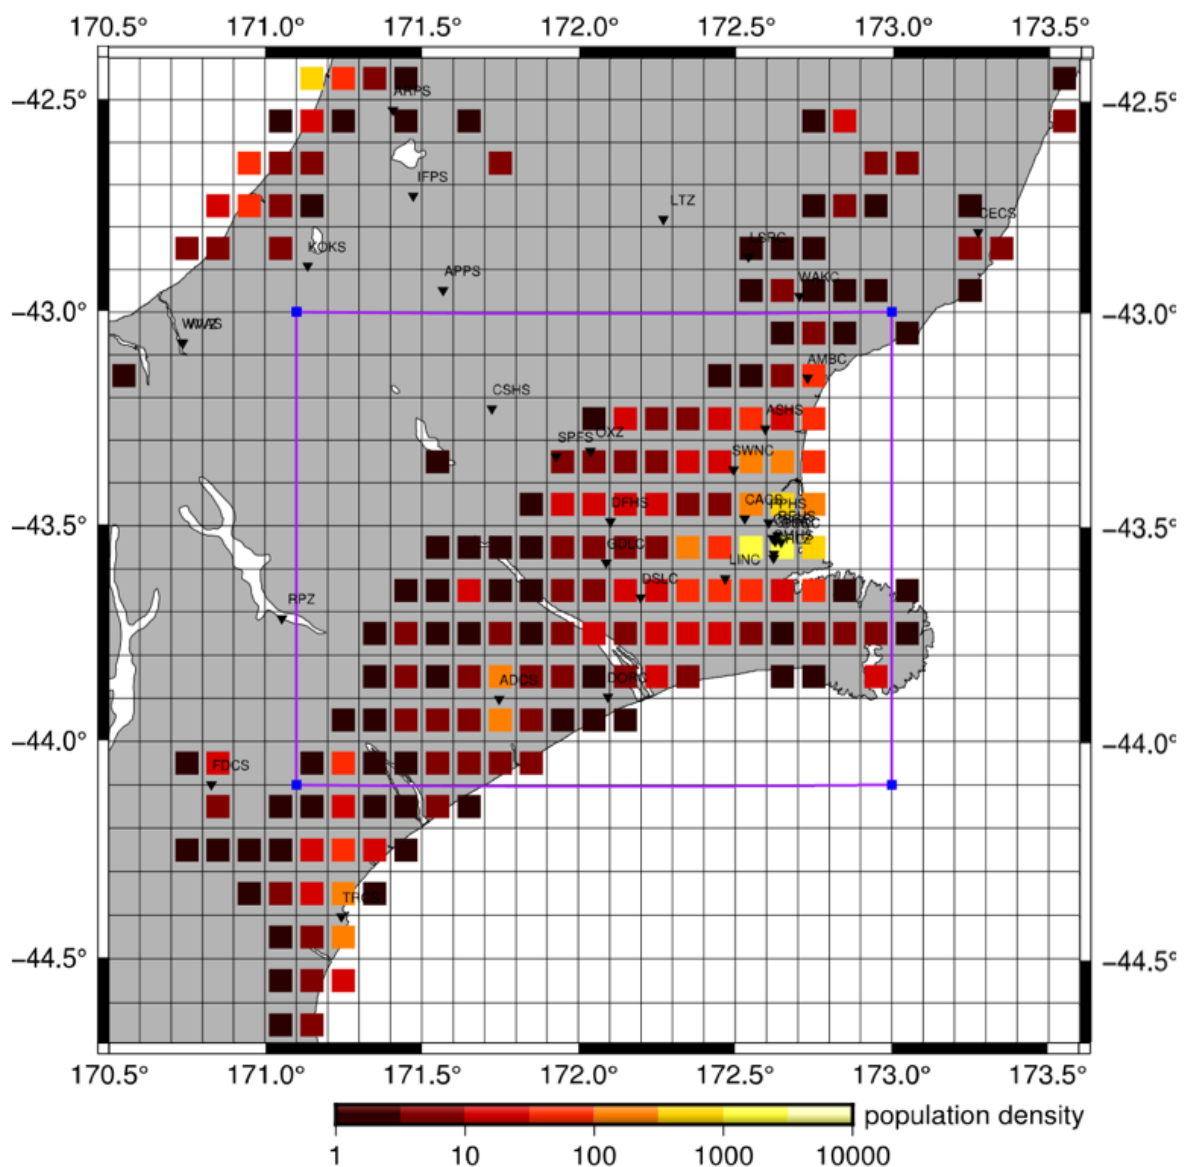


Figure 3-22 Population density of Canterbury region with average population density of each cell shown on a logarithmic colour scale. Gray cells have a population density under 1 person per data-cell (approximately 100 persons per grid cell).

The unweighted and weighted performance measures for the four different warning radii are shown in Table 3-12. At 30 km warning radius, the population-weighted correct timely alerts is 36% higher than the unweighted measure, but the population-weighted false alerts is 234% higher than the unweighted measure. This is consistent with the high fraction of correct timely alerts in the Christchurch region (see Figure 3-16) and the high concentration of the false alerts also in the Christchurch region (see Figure 3-17). For the larger warning radii, the population-weighted correct timely alerts is up to 56%. The population-weighted false alerts is also high, but not as dramatically higher than the unweighted measure as for the 30 km warning radius. With the population-weighted measures, it is still the case that the 40 km warning radius seems to give a significant improvement over the 30 km warning radius, but little improvement when the radius increases further to 60 km.

*Table 3-12 Population-weighted performance measures compared to the unweighted average performance measures.*

Warning radius	Average CTA <sub>norm</sub>	Population weighted CTA <sub>norm</sub>	Change	Average FA <sub>norm</sub>	Population weighted FA <sub>norm</sub>	Change
30 km	29%	40%	+36%	5.6%	19%	234%
40 km	32%	55%	+71%	9.8%	25%	159%
50 km	33%	56%	+67%	14.4%	32%	124%
60 km	33%	56%	+67%	19.8%	40%	99%

We also recalculated the best location to add a new station based on the population-weighted performance measures. Rather than adding a station in the far south-west of the central region (lat:-43.85, long: 171.45, see Figure 3-21), taking population density into account identified a location in the north-east, just north of Christchurch (lat:-43.65, long: 172.95). This is not surprising since it is clear that the best location to include a new station is likely to be around a densely populated area even though there are lots of other stations around.

### 3.5 Conclusions and limitations

This chapter discussed preparatory analyses that led to the more substantial analyses reported in Chapters 4 and 5. The overall conclusions and limitations of these preliminary

analyses are discussed in this section. The identified limitations are addressed in the design of the analyses reported in Chapters 4 and 5.

The GeoNet sensor latency experiments in Section 3.1 showed that it is possible to obtain most of the strong motion sensor data within two seconds of latency from the sensors if some modifications are applied to data acquisition devices to reduce the packet size. Therefore, it is reasonable to evaluate the PLUM EEW algorithm considering two second data latency. It may be possible to reduce the latency further by pre-processing done at the data acquisition devices of the sensor stations and sending the required information only (i.e., an “above the threshold” signal), and future NZ EEW system designers could consider that.

According to the analyses of Section 3.3, the EXSIM simulator is a reasonable ground motion simulator which could produce reasonable seismograms similar to the actual recorded data within the fault model and simulation uncertainties. Therefore, it is acceptable to use EXSIM to generate synthetic waveforms for the analyses of the thesis.

The Canterbury region PLUM analysis discussed in Section 3.4, which used 209 synthetic earthquakes, showed that a collection of synthetic earthquakes can be effectively used to measure the performance of the PLUM algorithm. That analysis suggested that a 30 km warning radius, as used in Japan, is not large enough for Canterbury. A 50 km or 60 km warning radius gives more correct timely alerts, but also gives a much higher number of false alerts. The results suggest 40 km warning radius might be a reasonable compromise. Unsurprisingly, the performance of PLUM was much better in the parts of the Canterbury region with a higher density of stations, and was very poor in the western part of the region where there is a very low density of stations. The analysis involving the addition of new stations showed that it is possible to use the simulated PLUM performance to determine appropriate locations to add stations. It also suggested that the new station locations should be identified by considering the EEW performance metrics if it is going to expand the existing sensor network for EEW purposes.

The chapter also explored a set of performance measures for PLUM and ways of presenting the outcome of the analyses, but also identified limitations of the measures and the methods.

The limitations of the analyses discussed in this chapter are presented below.

The set of 209 M6.5 earthquakes in the Canterbury PLUM analysis has an unrealistic distribution as they all have the same magnitude and have a uniform geographical distribution. This does not give a reliable prediction of the performance of PLUM on real earthquakes over the near future (next 100 years).

Even though the ground motion from the Darfield earthquake was simulated with realistic parameters, some parameters of the synthetic earthquakes are not realistic for the region, e.g., a strike of 0 degrees.

The list of sensor stations used in the analysis were obtained from GeoNet strong motion data tool but do not include all the stations that are now available. The 33 stations used were the ones with high-quality sensor data for the Darfield earthquake (GNS Science. (2022) and the analyses assumed a sensor infrastructure similar to that during the time of the Darfield earthquake. Further, GeoNet is constantly adding new stations and removing some old stations. Therefore, it is necessary to use a more up-to-date list of the current sensor stations.

There are slightly different versions of the PLUM algorithm that were tested in Japan and the US West Coast (Kilb et al., 2021). A simple version of the PLUM algorithm (single station detection) was used in these preparatory analyses to understand the PLUM operation and to design the framework to test it. A six consecutive samples window is used to reduce noise/error, but there is no confirmation with additional stations, so that the algorithm would be affected by noise from the sensors. Therefore, it has an optimistic prediction of the performance of PLUM, assuming less or no noise signals.

The detection of shaking was based on an approximation to the Japanese window method (Kodera et al., 2018; Yamamoto et al., 2008) which does not correspond well to the MMI measures used in NZ. The simulation also did not take into account the amplification factors for the site class and simply assumed that all the grid points were site class D.

The building blocks of the performance metrics were discussed in Section 3.4.1, but the performance metrics are not as informative as they should be. Firstly, the metrics do not distinguish two kinds of false alerts. Where an alert was received but there was no noticeable shaking, the alert is a genuine false alarm that will annoy end-users and is very undesirable for a real EEW system. Where an alert was received and shaking occurred which was noticeable but not strong, then the end users are much less likely to be annoyed, and such

alerts are much less of a problem for an EEW system. These latter alerts would be better classified as “precautionary alerts”. Treating precautionary alerts as false alarms in the performance measures is overly strict (Saunders et al., 2022a) and does not provide a useful measure for the effectiveness of the EEW algorithm. According to a survey by Goltz et al. (2024) regarding U.S. West Coast EEW, there was limited user preference for restricting alerts only to instances of damaging shaking.

Further, the normalisation of False Alerts did not produce a very informative measure. The number of false alerts was normalised against the total number of no-shaking events (FA + NANS). A grid point is labelled as NANS if an earthquake happened within or around the region, but the grid point receives no shaking above the threshold, and does not receive any alert. This means that  $FA_{norm}$  represents the times that an alert was received as a fraction of the times an alert should not have been received. However, the number of times an alert should not have been received is only well defined in a simulation where the number of earthquakes is well defined. It would be more appropriate to normalise FA against the total number of alerts (all types of alerts) received by a grid location.

In this chapter, missed alerts were categorised into two types. However, the distinction between missed alerts due to all the stations being too distant and ordinary missed alerts was not helpful.

Finally, some of the ways that population density was taken into account led to measures that were hard to interpret, particularly the population-weighted  $CTAnorm$  and  $FAnorm$ . It would be better to weight by population of the grid cells rather than population density.

Even though this Chapter didn't present any strong conclusions about PLUM algorithm performance in NZ, the exercises were very useful in designing the more useful analyses described in Chapters 4 and 5.

# Chapter 4 Evaluating PLUM on the Canterbury Region

## 4.1 Introduction

EEW systems rely on the input data coming from ground motion sensors. Therefore, earthquake ground motion waveform data is required to test the EEW algorithms. The best way is to use the data recorded by seismographs during past earthquakes. Although there is recorded data in New Zealand, it is insufficient for comprehensive analysis. Thus, we used the synthetic earthquake grid catalogue in the previous chapter and simulated synthetic seismograms to conduct a preparatory analysis to identify the performance of the PLUM algorithm to issue EEW for the Canterbury region. However, synthetic earthquakes regularly distributed along a grid pattern was unrealistic.

This chapter uses a more realistic synthetic earthquake catalogue—the RNC2 RSQSim catalogue (Shaw et al., 2022)—to mimic actual earthquake occurrences and to simulate synthetic seismograms for the selected earthquakes. That catalogue was generated using the RSQSim earthquake simulator (Richards-Dinger, 2012) for active faults in New Zealand. The total length of the catalogue is 220,000 years. Section 2.6 provided an introduction to this catalogue. The primary goal of the analyses discussed in this chapter is to explore the effectiveness of the PLUM algorithm for the Canterbury region over the next 100 years for significant earthquakes. The performance indicators are a modification of those used in the previous chapter. We still use correct timely alerts, missed or late alerts, and warning lead times, but we distinguish between precautionary alerts and false alerts and normalises the false alert measure in a more informative way. We also analyse performance considering population more effectively.

The analyses take into account the specific site classes at the stations and grid points when simulating the shaking and use a more complex version of PLUM that has more appropriate methods for detecting shaking and mitigating noise and errors in the seismograms than the method in the previous chapter. Section 4.1.1 describes the noise mitigation methods for the

PLUM algorithm used in Japan and the West Coast of the USA. The analyses also use a complete list of the GeoNet sensor stations that are now available in the Canterbury region.

The other objectives of the analyses are to determine appropriate parameters for PLUM, and study the effect of station outages. Chapter 5 will report on a similar analysis for the Wellington region.

#### **4.1.1 Noise mitigation for PLUM in Japan and USA**

PLUM is an efficient algorithm designed for rapid earthquake shaking detection by identifying strong shaking at sensor stations. However, when using real data, consideration must be given to reducing false detections caused by noise and errors by removing noise where possible and mitigating noise and errors that cannot be removed. Previous noise and error mitigation methods used for PLUM EEW are described in Kodera et al. (2018) for Japan and Cochran et al. (2019) for Southern California. Similar methods would need to be used for a NZ PLUM EEW system.

However, the analyses in this thesis use simulated data. The simulation data does not contain any shaking signals from sources other than earthquakes, nor does it contain instrumentation errors. Our main objective is to analyse the PLUM EEW performance and applicability considering the earthquake likelihood over the next 100 years. Therefore, the thesis does not explicitly focus on how to remove/mitigate the noise and errors of real data in an operational EEW system, although this would be an interesting and important area for future work. However, the simulation and analysis must take into account steps that would be taken to eliminate false detections.

The first step in both the Japan and the Southern California systems is to filter the data to remove noise, especially high-frequency and low-frequency noise. Kodera et al. (2018) use a bandpass filter adapted from the Japan Meteorological Agency (JMA) intensity filter, as shown in Figure 4-1.

The Japan JMA filter was originally designed to determine the intensity of earthquakes, and Kodera et al. (2018) used it for the PLUM EEW. The dominant frequency of that filter is 0.5 Hz. According to Hoshiya et al. (2010), the central frequency of 0.5 Hz is selected to characterise damaging strong motion in wooden frame houses. The filter decays slowly down

to around 0.05 Hz and up to 10 Hz. After reaching 10 Hz, it decays rapidly. This means it mainly removes the high-frequency component.

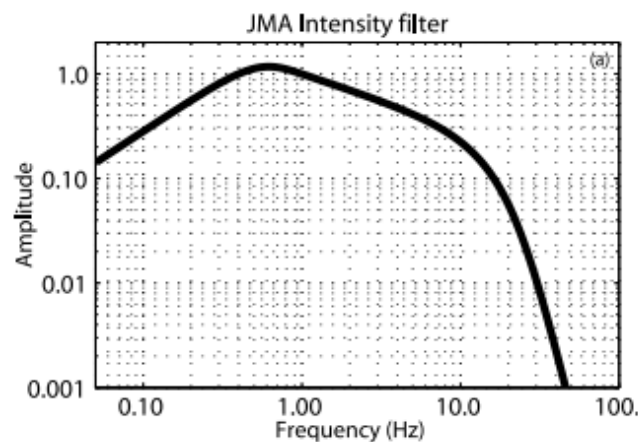


Figure 4-1 JMA intensity filter (from Yamamoto et al., 2008).

Cochran et al. (2019) use a 0.5–25 Hz bandpass filter for the Southern California data. This is similar to the JMA filter in that it is also band-pass, but it is shifted to a higher frequency than the JMA filter, with an upper boundary at 25 Hz rather than 10 Hz. The exact shape of the filter is not specified in the paper. A similar bandpass filter would need to be applied to the simulated shaking data used in this thesis.

The second step in both the noise and error mitigation approaches is to ensure that the event detection is not based on just a single data sample from a single sensor but seeks confirmation from multiple data samples. The Japan approach uses data across a “window” of data samples and the Southern California approach uses data from multiple sensors. In Japan, the Japan Meteorological Agency (JMA) employs a ground-motion intensity measure known as  $I_{JMA}$  (Yamamoto et al., 2008). The Japan PLUM system also uses real-time intensity  $I_R$ , which is similar to the  $I_{JMA}$  approach. According to that method, a signal is classified as an earthquake if a minimum of 0.3 seconds of samples (i.e., 30 samples at a 100 Hz sample frequency) are above the  $I_R$  threshold within an observation window of 5 seconds (Yamamoto et al., 2008; Kodera et al., 2018). They also incorporate this  $I_R$  value in the alert dissemination information.

In contrast to Japan, the intensity measure  $I_{MMI}$  is used in the USA (Wald et al., 1999) and New Zealand (Moratalla et al., 2021). Cochran et al. (2019) employed Peak Ground Acceleration (PGA) values to determine the Modified Mercalli Intensity (MMI) without calculating  $I_R$  values. Their approach differs from the Japanese method as it uses the peak value rather than 0.3

seconds of data values above across a predefined time window for data analysis. To mitigate the impact of noise spikes in real data, Kilb et al. (2021) suggested a dual-sensor strategy, utilising co-located accelerometers and velocimeters.

First, in each one-second time frame, they determine the maximum acceleration value (PGA) from the accelerometer and the maximum velocity value (PGV) from the velocimeter.

From PGA and PGV, they compute two values for MMI ( $I_{A:MMI}$  and  $I_{V:MMI}$ ) using the PGM to MMI conversion equations of Worden et al. (2012). Then, they examine the intensity ratio of the two MMI values ( $I_{A:MMI} / I_{V:MMI}$ ) and assign the  $I_{MMI}$  as described in Table 4-1 below.

*Table 4-1 Conversion of  $I_{A:MMI}$  and  $I_{V:MMI}$  to  $I_{MMI}$*

	Intensity Ratio $I_{A:MMI} / I_{V:MMI}$	$I_{V:MMI}$	Assumption	Final $I_{MMI}$ value
Case A	0.9 - 1.1	--	Ratio is within acceptable range	$( I_{A:MMI} + I_{V:MMI} )/2$
Case B	$\leq 0.9$ or $\geq 1.1$	$I_{V:MMI} < 3$	Ratio is not within acceptable range. Acceleration data suspect.	$I_{V:MMI}$
Case C	$\leq 0.9$ or $\geq 1.1$	$I_{V:MMI} \geq 3$	Ratio is not within acceptable range. Velocity data may be clipped.	$I_{A:MMI}$

If a station does not have both sensors, they derive velocity data from the accelerometer or acceleration data from the velocimeter by analysing a one-second data frame. They then perform the same operation as described above. However, unlike the Japanese approach, they did not rely on a 0.3-second window criterion and used single PGA and PGV values to derive the MMI.

The simulated data in this thesis has reliable acceleration data, and therefore the analyses can simply use the PGA value without needing velocity data to mitigate noise and errors or having to use multiple data values over a window. At worst, this may slightly overestimate the performance of the PLUM algorithm since the noise mitigation on real data would not be perfect.

A third noise mitigation strategy is to require confirmation of an event from another station. Southern California used a dual-station approach for earthquake event declaration, necessitating confirmation from a neighbouring station (Cochran et al., 2019; Kilb et al., 2021). This method requires the detection of an earthquake by at least two stations within a defined range to issue alerts. Cochran et al. (2019) achieved better detection performance when at least two stations were located within 60 km. Kilb et al. (2021) consider a confirmation within 15 seconds (i.e., from a sensor within 45 km, assuming S-wave velocity as 3 km/s). The initial station employs one  $I_{MMI}$  threshold (e.g.,  $I_{MMI}$  4.0, as tested in Southern California) to identify the event, but a neighbouring station can confirm this detection with a lower  $I_{MMI}$  threshold (Cochran et al. recommended  $I_{MMI}$  2.5 for Southern California).

The analyses in the thesis use a similar confirmation strategy to the one used in the Southern California PLUM testing (Cochran et al., 2019; Kilb et al., 2021).

## 4.2 Methodology

This section outlines the steps to analyse the performance of the PLUM algorithm for EEW in the Canterbury region. There were five key steps, which are explained in the following subsections.

1. A region of interest (Canterbury) was selected, target grid locations in the region were defined and the locations of all the relevant GeoNet sensors were identified. The site classifications for those grid and sensor locations were identified to use for the simulation.
2. A collection of earthquakes from the synthetic catalogue was selected. One hundred sets of random 100-year-long periods from the 220000-year catalogue were selected, and all the earthquakes in those periods with magnitude greater than 6.0 were identified under the assumption that earthquakes with magnitudes less than 6.0 are unlikely to cause significant damage. Each set represents a possible scenario of strong earthquake events for the next 100 years in New Zealand. Then, the sets were refined to include only the earthquakes that could impact the chosen region of Canterbury.
3. Synthetic seismograms for all selected events were generated using the EXSIM ground motion simulator for all sensor (GeoNet strong motion) and grid locations.

4. The PLUM algorithm was used on the generated seismograms using an appropriate noise and error mitigation strategy.
5. The expected EEW performance of the algorithm over each 100-year period was analysed in terms of alert correctness and timeliness, first ignoring population density, then taking the population distribution into account.
6. Finally, the expected performance was re-analysed under two different failure scenarios.

#### **4.2.1 Region of interest**

The Canterbury region was selected for the preliminary analyses in Chapter 3 as it has a richer collection of actual datasets available than other regions, especially from the Canterbury earthquake sequence. The same region is used to test EEW with the RNC2 synthetic catalogue in this chapter.

The region of Canterbury where the EEW algorithm is evaluated is shown in Figure 4-2 by a red rectangle ( $-44.10 \leq \text{lat} \leq -43.00$ ,  $171.10 \leq \text{long} \leq 173.00$ ), along with the grid cells (containing at least some part on land) used in the analysis.

Figure 4-3 shows the larger area, containing the selected region, and indicates the 131 GeoNet strong motion sensor stations (triangles) that were selected to test the algorithm for that region. The 131 sensors are located within 250 km of the centre of the selected region. That range of sensor stations was selected considering the appropriate coverage of sensors according to the PLUM warning radius ranges of 30-60 km. This includes all the sensors that might be involved in creating or confirming alerts for the Canterbury region, however, excludes the large number of sensors of the Wellington and Dunedin regions, which could not affect the operation of the PLUM algorithm for the Canterbury region. As in Section 3.4, 209 grid points ( $0.1^\circ \times 0.1^\circ$ ) are defined in the inner rectangular region (the red rectangle in Figure 4-2), of which 182 grid points on land are selected as the target locations for analysing the EEW performance.




Figure 4-2 The selected Canterbury region with  $0.1^\circ \times 0.1^\circ$  grid. Region is the area inside the red rectangle. Grid cells containing red dots are the locations where alerts are received and analysed. The urban areas of Christchurch and Ashburton are labeled.




Figure 4-3 Canterbury region with sensors. Red rectangle shows region containing all the grid locations at which the PLUM algorithm is to be evaluated. Black triangles indicate all the GeoNet strong motion sensors (as of September 2023) that could possibly issue or confirm alerts to the central region.

#### 4.2.2 Site classification

Site classifications of grid cells and sensor stations are considered in this study since ground motion amplitudes at a location depend on the site class of the location.

The site class of each grid cell is determined from a New Zealand VS30 model (Perrin et al., 2015). The resolution of the VS30 model is  $0.01^\circ \times 0.01^\circ$  degrees, while the grid cell size is  $0.1^\circ \times 0.1^\circ$ . Therefore, the VS30 site class closest to the mid-point of the centre of the grid cell is considered to be the site class of the grid cell. Effectively, this is pretending that a sensor is installed close to the mid-point of the grid cell, just as the Japan PLUM implementation has

basic seismic intensity-measuring sensors installed at each target location as observation stations in addition to the strong motion seismometers.

The following VS30 value ranges (NEHRP, 2023) are used to determine the site classes from the VS30 model.

$VS30 > 1500 \text{ m/s}$	→	Class A (Hard Rock):
$1500 \geq VS30 > 760 \text{ m/s}$	→	Class B (Firm to Hard Rock)
$760 \geq VS30 > 360 \text{ m/s}$	→	Class C (Dense Soil and Soft Rock):
$360 \geq VS30 \geq 180 \text{ m/s}$	→	Class D (Stiff Soil):
$180 > VS30 > 0$	→	Class E (Soft Soil):

If the VS30 value is less than or equal to 0, the site is classified as a water site and is excluded.

If the mid-point of a grid cell is on water (or the grid cell is partially covered by water), the site class is defined manually by considering the VS30 site classes of the points in the grid cell that are on land. Mostly, these grid cells are classified as site class D.

A total of 131 stations within a 250 km radius of the region's centre were used for this study (GNS Science, 2023b). The site classification of most of these stations was obtained from Kaiser et al. (2017b). However, site classification data is unavailable for some of the stations. These include recently installed stations and sensors installed within buildings. For these stations, we assumed a site class of B, as it is the hardest site class among all other stations and therefore is a conservative choice because the shaking intensity is less likely to be above the threshold than if a softer site class were chosen.


#### **4.2.3 Random selection of sub-catalogues from the main catalogue**

The version of the catalogue that we are using spans over 220,000 years. However, the interest period concerning EEW system implementation will be relatively short. New Zealand Building Code requires that all buildings be designed and constructed to be durable enough to last at least 50 years. Therefore, the period of interest for these EEW algorithm analyses should be at least 50 years. The period of interest at most will be 100 years, as it is not reasonable to expect the current technology to be applicable for more than 100 years. Therefore, we need samples of earthquakes that represent what might happen over the next

100 years. We randomly selected 100 sub-catalogues, each of 100 years duration, from the full catalogue to get a statistically reliable prediction of what might happen over the next 100 years.

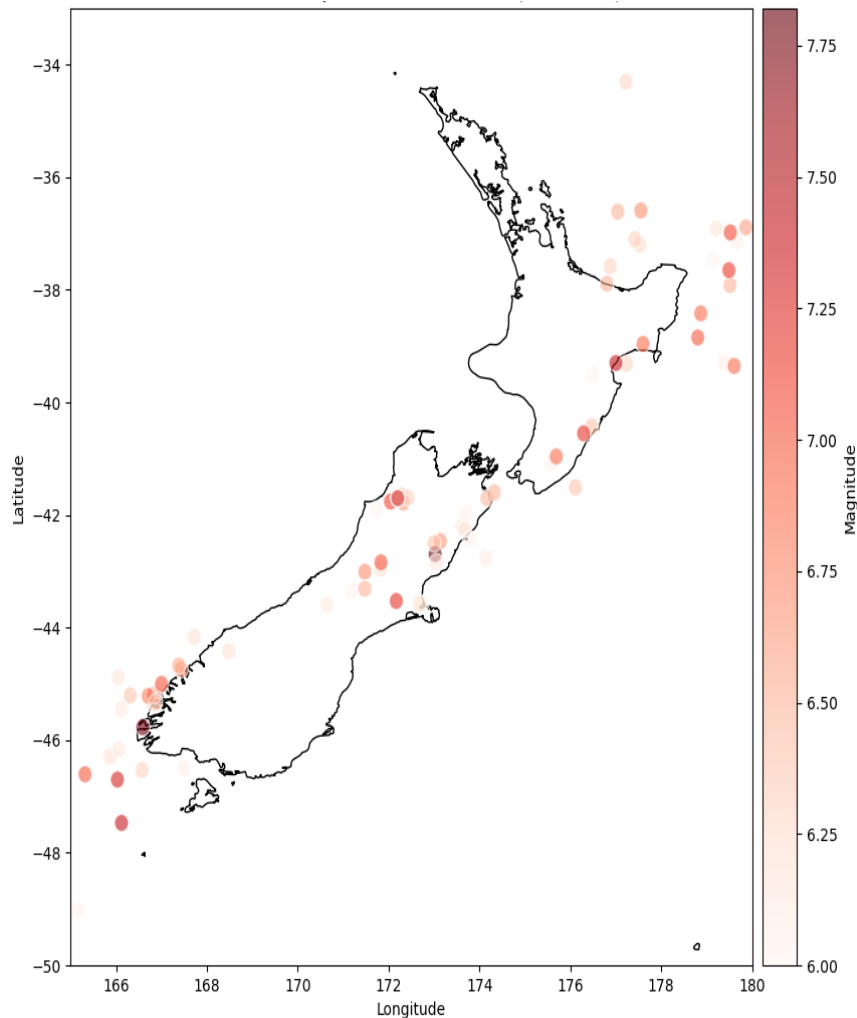
We focused on earthquakes with magnitudes greater than 6. Smaller earthquakes can also cause damage and MMI levels above V, but these effects are usually restricted to areas near the epicentre, where early warnings are impossible.

Earthquakes from a selected 100 year-long sub-catalogue are illustrated in Figure 4-4 as red circles, along with the earthquakes from the other 99 sub-catalogues coloured grey.



*Figure 4-4 Synthetic earthquakes: all earthquakes greater than magnitude 6 in the 100 sets of 100-year long sub-catalogues. Red circles indicate earthquakes from one selected 100-year long sub-catalogue and grey circles indicate earthquakes from the other sub-catalogues. Darker colour circles indicate higher magnitude.*

The New Zealand historical catalogue of earthquakes (GeoNet data) with magnitude greater than 6.0 and depth less than 30 km (since the RSQSim model only simulated shallow events) is illustrated in Figure 4-5 and we can see the synthetic earthquakes in the sub-catalogues are spatially consistent with this data.



*Figure 4-5 Historical catalogue of shallow earthquakes (M6.0+) in the last 100 years: 1924-2023.*

The initially selected earthquakes were further refined to select the earthquakes that could produce shaking equal to or above MMI V at some location within the selected Canterbury region—earthquakes that are too far away from the region to cause damage can be removed. To select the potentially damaging events, we first included all earthquakes of magnitude 6.0 and above that have their centroid within the selected region of Canterbury. To identify earthquakes outside the region, which could produce MMI V or above shaking at any part of the region, we calculated the shortest distance from each earthquake fault to the centre of

the region. For an earthquake to be included, this distance had to be below a threshold determined from the earthquake magnitude. Thresholds for the distance could be determined using the relevant ground motion prediction equation (GMPE). Instead, we used the NGA-West2 ground motion database tool (Pacific Earthquake Engineering Research Center, 2023) to calculate the distance over which specific earthquake magnitude ranges could produce shaking at MMI V or above. The NGA-West2 tool uses attenuation relationships modelled using a large set of ground motions recorded worldwide of shallow crustal earthquakes in active tectonic zones (Ancheta et al., 2013; Abrahamson et al., 2014). These models are applicable to the Canterbury region. An additional 100 km was added to the thresholds to account for the distance from the centre of the region to a location at the edge of the region.

Based on the GMPE data, locations at the following distances from the fault might experience shaking at MMI V or above:

$6 \leq \text{Magnitude} < 7$ : within 140 km

$7 \leq \text{Magnitude} < 8$ : within 240 km

$8 \leq \text{Magnitude} < 8.5$ : within 300 km

The following algorithm outlines the selection criteria for earthquakes to be included in the sub-catalogues based on earthquake magnitude (M) and proximity of the fault to the centre of the region:

If  $M < 6$ :

Ignore

Else if centroid within the region:

Select

Else if shortest distance to fault  $\leq (140 + 100)$  km:

Select

Else if  $M > 7$  AND shortest distance to fault  $\leq (240 + 100)$  km:

Select

Else if  $M > 8$  AND shortest distance to fault  $\leq (300 \text{ km} + 100)$  km:

Select

Else if  $M > 8.5$ :

Select

Else:

Ignore

Note that the thresholds are set generously so the selection criteria will include some earthquakes that do not generate strong shaking in the Canterbury region. These earthquakes will require simulation time, but if the shaking is below the strong shaking level, the PLUM algorithm will effectively ignore them, so they will not affect the analysis.

946 earthquakes were selected to construct the 100 sub-catalogues most of which are crustal earthquakes. Since subduction earthquakes need to be simulated using a different set of parameters, we examined the locations of all earthquakes greater than magnitude 7.5. Based on their location and magnitude, 13 earthquakes were identified as subduction earthquakes, all of which are greater than M8.5. One very distant magnitude 8.9 subduction earthquake from those 13 earthquakes could not be simulated because the simulation software could not handle such a large earthquake over such large distances. We, therefore, ignored that earthquake. The maximum PGA experienced by the region from two similar (but closer) earthquakes of magnitude 9.0 and 8.9, is lower than the MMI V threshold, and therefore, the ignored earthquake would probably not generate shaking above the threshold to the region.

The final result was that 945 earthquakes were selected for simulation, with 933 being crustal and 12 being subduction. The average number of earthquakes in one 100-year period was 9.45, but there was considerable variability. Figure 4-6 is a histogram showing the distribution




Figure 4-6 Histogram showing distribution of number of earthquakes: each bar shows the number of periods that had a given number of earthquakes.

of the number of earthquakes per 100-year period. Except for one period, all had at least 5 earthquakes. The largest number of earthquakes in a period was 19. The majority of periods (68 periods) experienced between 7 to 11 earthquakes. The distribution is roughly Poissonian, with a peak at around 8 to 9 earthquakes per period, indicating that this was the most common number of earthquakes per period in the dataset.

Figure 4-7 shows the centroids of all the 945 selected earthquakes of all the 100 random 100-year periods, and earthquakes of a single (randomly chosen) 100-year period are highlighted by red circles. According to Figure 4-7, there are no earthquakes immediately around or under Christchurch. This is due to the catalogue's limitation, as it did not include the blind fault under the city. This suggests that the catalogue might underestimate the occurrence of earthquakes near Christchurch City. Interpreting the results needs to take this limitation into consideration.

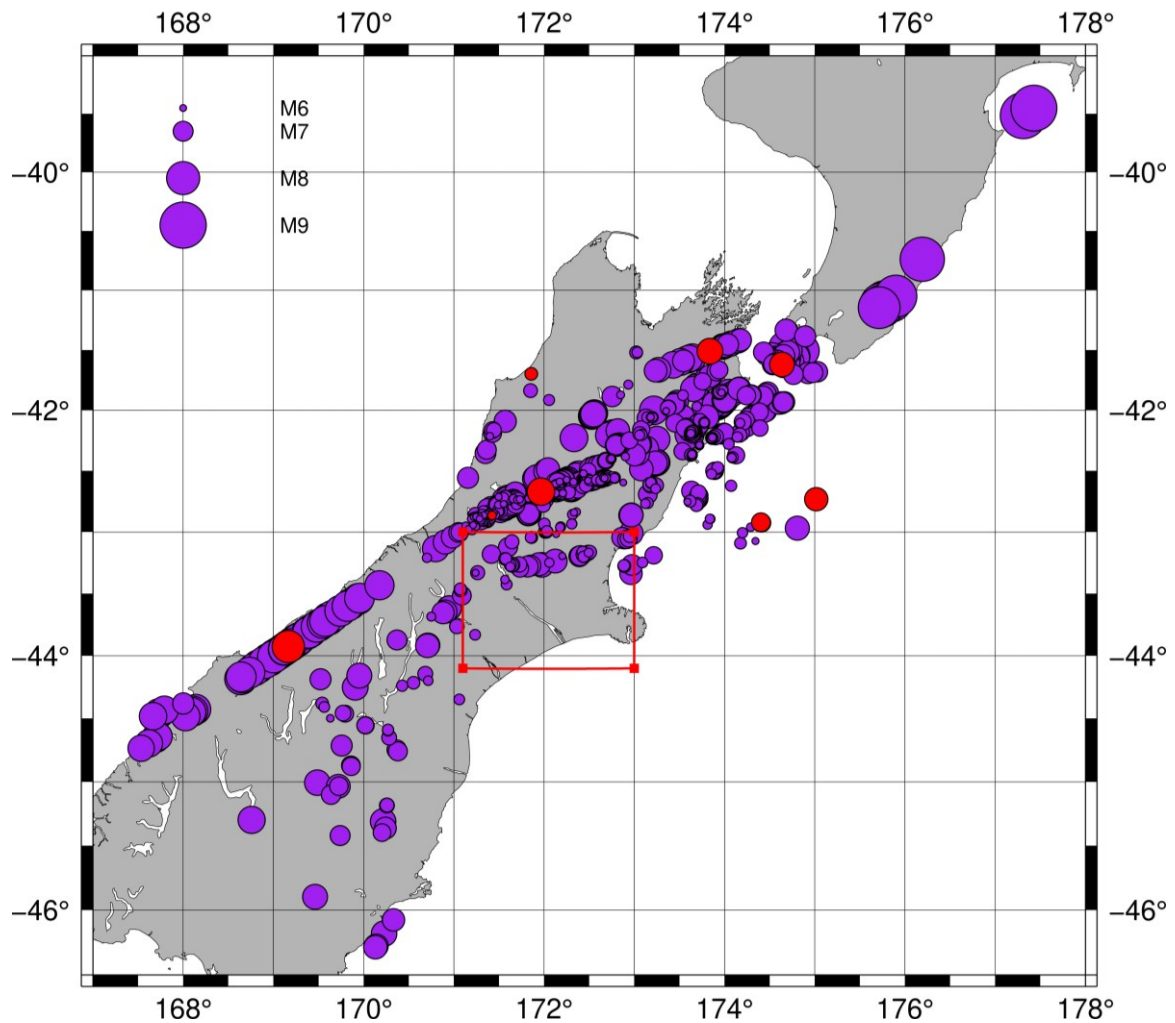


Figure 4-7 All earthquakes that could affect the Canterbury region from the random 100 sets of the 100-year periods from the 220000-year catalogue. The circles are the centroids of the earthquakes, and the size of the circles represents the magnitude. The red circles are earthquakes from one selected 100-year period (#69). The red rectangle is the inner region containing 182 grid points ( $0.1^\circ \times 0.1^\circ$ ) where performance of the EEW algorithm is evaluated.

#### 4.2.4 Synthetic seismogram generation for the catalogue events

Testing the PLUM EEW algorithm requires earthquake waveform data for each event at the locations of all the sensor stations and at the centre of each  $0.1^\circ \times 0.1^\circ$  grid cell. The EXSIM ground motion simulator was used to generate these synthetic seismograms. The general EXSIM input parameters were based on the analysis in chapter 3 where EXSIM was validated to the recorded waveform data for the Darfield earthquake.

As described in previous chapters, EXSIM requires key characteristics of each event, including the fault position. However, many of the catalogue events comprise multiple faults. EXSIM is designed to simulate a single fault rupture. Although it can be extended to multiple fault

rupture under certain limitations, obtaining and using the individual fault data for multiple faults would be very complex. Therefore, to keep the simulation simple and feasible, we modelled and simulated all the events as single faults and computed fault positions from the catalogue data.

All earthquakes were separated into two major types, crustal and subduction, which were simulated separately using the relevant characteristics as described below.

### **Crustal earthquakes**

For the crustal earthquakes, the dip was assumed to be 90° and the depth of the top of the fault was assumed to be 1 km. The fault type of the events was assumed as strike-slip. The slip values were determined by the simulator randomly. S-wave velocity was considered to be 3.7 km/s (see Section 3.3.1) and the ratio of rupture velocity to S-wave velocity was considered to be 0.8 for the selected region. According to Holden & Kaiser (2016) and the analyses in Section 3.3, a stress drop of 9 MPa was assigned for the earthquakes.

EXSIM also requires the location of the corner on the upper edge of the fault, but the catalogue did not provide this. To work this out, we assumed the centroid was at the middle of the fault and used the fault bearing angle (strike) and fault length to compute the corner location. The fault strike was determined from the direction from the centroid to the hypocentre of the fault (both locations are specified in the catalogue). The fault length was determined from the magnitude using the Wells and Coppersmith formula (Wells and Coppersmith, 1994) as shown in equation 4.1, and both the calculated corner on the upper edge of the fault and the fault bearing angle (strike) were given to the EXSIM simulator. Note that EXSIM also uses the same Wells and Coppersmith formula to determine the fault dimensions from the magnitude itself.

$$\text{fault length (km)} = e^{(-2.57 + 0.62M)} \quad (4.1)$$

where  $M$  is the magnitude of the earthquake.

According to Shaw (2023), there are some imperfect approximations in Wells and Coppersmith's formulas. We used Wells and Coppersmith to estimate the fault length from the magnitude in part because the EXSIM simulator uses Wells and Coppersmith internally, and it would be problematic to use Wells and Coppersmith for one part of the analysis and a

different estimator for other parts of the analysis. Also, we think that Wells and Coppersmith's estimation is adequate in this case because the catalogue data consists of hypothetical earthquakes with insufficient data to perform more accurate estimations.

Initially, the simulated hypocentres were determined randomly by the simulator, but we found an issue with the random numbers in the simulation process, so we fixed the hypocentre at the same location as the catalogue centroid, i.e., in the middle of the fault.

### **Subduction earthquakes**

The selected earthquakes include several subduction events from the Hikurangi subduction zone, which could impact the Canterbury region. These subduction earthquakes were simulated under the assumption that the fault had a fixed strike angle of 210°, a dip angle of 10°, and a depth of 1 km (Stirling et al., 2012). The strike angle of 210° represents an average value derived from Stirling et al. (2012), which reported strike angles between 200° and 220°. The events were considered reverse slip earthquakes, and the Wells and Coppersmith's parameters were used to calculate the fault length and fault width, as shown in Eq. 4.2:

$$\text{Fault length (km)} = e^{(-2.42 + 0.58M)} \quad \text{Fault width (km)} = e^{(-1.61 + 0.41M)} \quad (4.2)$$

where  $M$  is the magnitude.


The location of the top right corner of the fault was calculated from the centroid, the computed fault length and width, and the assumed dip angle of 10°. Other parameters are identical to the previous crustal events.

For each earthquake, EXSIM was run separately to obtain the waveform for each sensor location and each grid cell centre.

#### **4.2.5 Noise and error mitigation**

This research used the prior work on PLUM EEW by Kodera et al. (2018) in Japan and Cochran et al. (2019) in Southern California to determine appropriate noise and error mitigation measures as discussed in this section. Both of those prior works are based on real data, whereas our analyses use simulation data. The simulation data does not contain seismic signals from sources other than earthquakes. It also does not contain instrumentation

glitches. Since our main objective is to analyse the PLUM EEW performance and applicability considering the earthquake likelihood over the next 100 years, this thesis does not explicitly focus on the noise and errors of real data in an operational EEW system, which will be an interesting area for future work. However, for the simulation and analysis to be applicable, it must take into account the necessary steps that would be taken to eliminate false detections. First, it is necessary to filter the data to remove noise, especially high-frequency and low-frequency noise. Section 4.1.2 described the bandpass filters used by Kodera et al. (2018) in Japan and Cochran et al. (2019) in Southern California. For our EXSIM simulation, we use the bandpass filter profile shown in Figure 4-8 Figure 4-8 to remove low and high-frequency noise. This filter is broader than the filter used in Japan. It includes the dominant frequency (0.5 Hz) of the Japanese filter but is narrower than the filter used in Southern California (0.5-25 Hz), especially at the higher frequencies. Therefore, it is consistent with the range of filters used in the literature. It is appropriate to use a broader filter than used in Japan because our interest is not only in scenarios like wooden frame houses in Japan but also in a broader range of situations. Note that the same filter was used for our validation and calibration analyses against real data from the Darfield earthquake.



*Figure 4-8 Filter applied in the EXSIM simulation*

The second step is to use a confirmation strategy to ensure that alerts do not depend on a single sensor that might have errors due to noise or any other issue.

We used a two-station confirmation strategy like that used in Southern California (Cochran et al., 2019; Kilb et al., 2021) which requires confirmation of shaking (of at least MMI-IV) from a

nearby station within some confirmation window before issuing an alert. Kilb et al., (2021) use 45 km as the confirmation radius for the West coast, USA data, based on the sensor distribution there. The density of the GeoNet strong motion sensors in the Canterbury region is not high enough to support this radius. However, there are at least two stations within 80 km of each station in the Canterbury region and therefore we set the confirmation radius at 80 km. The two-station strategy requires confirmation of from another station within 80 km of the target grid location. It will eliminate any alerts that were due to noise assuming that no other earthquake event with an intensity (i.e., greater than MMI-IV) occurs in the region within the confirmation window, and that there are no additional noise detections in that period at other stations within 80 km. The time for an S-wave to travel 80 km will be 22 seconds, assuming an S-wave velocity of 3.7 km/s, so the confirmation window should extend at least 22 seconds each side of a detected strong shaking event. Unlike the Southern California strategy, which requires the confirmation to occur after the alerting station detects strong shaking, our strategy allows the confirmation to occur either before or after the strong shaking at the alerting station. Even though there are no noise events in the simulated waveforms, it is important to use the two-station method to obtain more realistic warning times and alert performances mimicking an actual system.

The confirmation window should be long enough (both before and after the strong shaking is detected at a station) to ensure that confirmation messages based on the same earthquake event have time to arrive at the station. However, making the confirmation window too long could allow false confirmations from unrelated events and would also increase the chance of confirmation from noise events. Choosing a good confirmation window would require analyses on real data. However, for the simulated data, there are no unrelated events and no noise, so it does not affect the results to allow an indefinite confirmation window.

The event detection procedure of these analyses did not use the Japanese strategy of seeking confirmation from multiple data samples across a “window”, as we do not use  $I_R$  relationships in New Zealand and we use PGA to MMI relationships (Moratalla et al, 2021) and using a window also delays the alerts. Also, because we cannot assume that the GeoNet sensor stations have dual acceleration and velocity sensors, we do not use the co-located sensor strategy used by Cochran et al. (2019). Instead, the algorithm identifies the first time that the simulated shaking waveform reaches a value higher than the threshold, which corresponds

to using the PGA value (peak ground acceleration). There are two separate thresholds, a lower threshold for issuing a confirmation message ( $\tau_l$ ), and a higher threshold for issuing an alert ( $\tau_h$ ).

The PLUM algorithm with second station confirmation is as follows (refer to Figure 4-9 Figure 4-9).

#### PLUM Algorithm with second station confirmation.

- When a station detects shaking (PGA at  $\tau_l$  or above):  
it sends a confirmation message to all grid points within the confirmation radius.
- When a station detects strong shaking (PGA at  $\tau_h$  or above):  
it sends an alert message to all grid points within the warning radius.
- When a grid point receives an alert message:
  - If it has already received a confirmation message (at a time within the confirmation window of the alert time), it will immediately issue the alert.
  - Else, it holds the alert message and waits for a confirmation message.
    - If it receives a confirmation message before the end of the confirmation window, it will then issue the alert.
    - If it does not receive a confirmation message within the confirmation window, it will abandon the alert message.




Figure 4-9 PLUM confirmation and warning. When grid point receives an alert from a station within the Warning Radius, it must also receive a confirmation from another station within the Confirmation Radius before processing the alert.

#### 4.2.6 PLUM algorithm

For each earthquake, after obtaining the waveform data for all station locations and all grid locations, as described in Section 4.2.4, the PLUM algorithm was evaluated on the waveform data to determine the outcome of the PLUM algorithm at each grid point. The possible outcomes at a grid point are given in Table 4-2. These are similar to the earlier analyses in chapter 3, except that the missed alerts due to no station being within the warning radius (MANS) have been combined into the other MA and the false alerts have been split into precautionary alerts (PA) and real false alerts (FA).

A precautionary alert is where an alert is received but the grid point only experiences weak shaking. A false alert is where an alert is received but any shaking at the grid point is below the level of weak shaking. The weak shaking means that people receiving a precautionary alert will recognise that the alert is from a real earthquake, but it was not strong enough shaking to cause damage (at least MMI IV, but less than MMI V). On the other hand, people receiving a real false alert will consider the system to be a fault since they will have no evidence of an earthquake, and therefore no justification for being alerted. The distinction is made because real false alerts are perceived by most end users to be a much worse problem in an EEW system than precautionary alerts, so that treating them as equally bad results in an unnecessarily harsh evaluation (Saunders et al., 2022a).

*Table 4-2 Possible outcomes of PLUM*

Outcome	Meaning
CTA:	Correct Timely Alert, where a confirmed alert at the grid point is processed before strong shaking arrives. The Warning Time is the time between the processing of the alert and the arrival of the strong shaking.
LA:	Late Alert, where the first confirmed alert at the grid point is processed after strong shaking arrives.
MA:	Missing Alert, where strong shaking arrives at the grid point, but no confirmed alert is received.
PA:	Precautionary Alert, where a confirmed alert arrives at the grid point, but the grid point only receives weak shaking.
FA:	False Alarm, where a confirmed alert is processed at a grid point which does not receive any noticeable shaking.
NANS:	No Alert, No Shaking

To measure the performance of the PLUM algorithm, it is simpler (and more efficient) to “invert” the algorithm in Section 4.2.5, focusing on the grid points, rather than on the sensors that issue the alerts. The algorithm used for this simulation of PLUM was as follows.

### PLUM Algorithm with confirmation, for analysis

#### Parameters:

- $\tau_h$  (strong shaking detection threshold)
- $\tau_l$  (weak shaking detection threshold)
- warning radius
- confirmation radius
- comm-time [communication latency]

#### Output:

- outcome and warning-time for each grid point for each earthquake.

#### Algorithm:

##### For each grid point:

compute list of stations within warning radius of grid point  
 compute list of stations within confirmation radius of grid point

##### For each earthquake:

##### For each station:

$ws\text{-time}[station] = \text{first time when } pga \geq \tau_l \text{ (infinity if no such time)}$   
 $ss\text{-time}[station] = \text{first time when } pga \geq \tau_h \text{ (infinity if no such time)}$

##### For each grid point:

$ws\text{-time}[grid\text{-pt}] = \text{first time when } pga \geq \tau_l \text{ (infinity if no such time)}$   
 $ss\text{-time}[grid\text{-pt}] = \text{first time when } pga \geq \tau_h \text{ (infinity if no such time)}$

[Determine whether a confirmed alert is processed:]

$alert\text{-time} = \text{infinity}$

$confirmation\text{-time} = \text{infinity}$

**For each** station in list of stations within warning radius of grid point

**If**  $ss\text{-time}[station] < alert\text{-time}$  **then**

$alert\text{-time} = ss\text{-time}[station]$

$alert\text{-station} = station$

**If**  $alert\text{-time} < \text{infinity}$  **then**

**For each** other-station within confirmation radius of grid point

**If** other-station  $\neq$  alert-station **and**

$ws\text{-time}[other\text{-station}] < confirmation\text{-time}$  **then**

*confirmation-time = ws-time[other-station].*  
*alert-time = max(alert-time, confirmation-time)*

[Determine outcome of event at grid point]:

**If** *ss-time[grid point] < infinity* **then** [*ie, strong shaking at grid point*]

**If** *alert-time = infinity* **then**

*outcome[earthquake, grid point] = MA*

**Else**

*warning-time[eq, grid point] = alert-time –*

*ss-time[grid point] – comm-time*

**If** *warning-time[eq, grid point] ≥ 0* **then**

*outcome[eq, grid point] = CTA*

**Else**

*outcome[eq, grid point] = LA*

**Else If** *ws-time[grid point] < infinity* **AND** *if alert-time < infinity* **then**

*outcome[eq, grid point] = PA*

**Else If** *alert-time < infinity* **then**

*outcome[eq, grid point] = FA*

**Else**

*outcome[eq, grid point] = NANS*

Implementing PLUM requires choosing values for the parameters. The parameter values for the Canterbury experiment are listed in Table 4-3.

Different ground motion thresholds have been suggested by researchers. In Japan, a public EEW alert is issued if the expected ground motion exceeds JMA Intensity 5L (approximately equivalent to MMI VII-VIII) (Kodera et al., 2020). The ShakeAlert EEW System of the US West Coast used MMI IV as the original threshold and later, delivery criteria have been refined (Cochran & Husker, 2019; Saunders et al., 2022b) to better match different user needs. Today, general public alerts delivered via Wireless Emergency Alert capable devices use a threshold of MMI IV+, while smartphone apps allow users to choose MMI III+ (ShakeAlert, 2024). On Android devices, push notifications are sent at MMI III–IV, and a full-screen takeover at MMI V. Automated systems such as signal light panels and sirens also could trigger at MMI III+, as do machine-to-machine alerts (ShakeAlert, 2024). Minson et al., (2019) stated that it is possible to use lower thresholds, especially for false-alert-tolerant users. Even though it could generate more false alerts, this will minimise missed alerts and improve the

usefulness of EEW. Therefore, the threshold can be lower or higher than the damage threshold according to the user preferences.

Becker et al. (2020b) found that New Zealanders generally prefer higher warning thresholds around MMI V-VI which aligns with preferences observed in Japan but exceeds the threshold used on the U.S. West Coast (Minson et al., 2019, 2022; Nakayachi et al., 2019). In the survey of Becker et al. (2020b), a majority of respondents indicated an alert at MMI V or above as their preference. Consequently, for our PLUM algorithm tests in the New Zealand context, we adopt MMI V as the detection and target threshold, which is consistent with public expectations and international practices.

*Table 4-3 PLUM parameters for Canterbury catalogue earthquakes experiment.*

$\tau_h$ (Detection threshold and Strong Shaking threshold)	42 cm/s <sup>2</sup> (PGA corresponding to MMI V, which is considered strong shaking, with the potential for damage)
Warning radius:	30 km, 40 km, 50 km, or 60 km
$\tau_l$ (Confirmation threshold and Weak Shaking threshold)	13 cm/s <sup>2</sup> (PGA corresponding to MMI IV, which is considered weak shaking. Shaking below MMI IV is not usually noticed.)
Confirmation radius	80 km
Confirmation window	Indefinite
Communication latency	2 sec

Initially, Wald et al., (1999) PGA to MMI relationships were used to define the detection and confirmation thresholds, which gave a detection threshold  $\tau_h$  of 38 cm/s<sup>2</sup>, but now we have more accurate relationship equations for New Zealand from Moratalla et al. (2021). According to Equation 4.3 (derived from Equation 3 and Table 2 of Moratalla et al., 2021), PGA thresholds can be calculated as below.

$$\begin{aligned} \text{Log (PGA)} &= (\text{MMI}-1.7601) / 1.992 && \text{if MMI} < 5.5277 \\ \text{Log (PGA)} &= (\text{MMI}+1.9095) / 3.9322 && \text{if MMI} \geq 5.5277 \end{aligned} \quad (4.3)$$

Therefore:

$$\tau_l \text{ for MMI IV} = 10^{(4-1.7601)/1.992} = 13 \text{ cm/s}^2$$

$$\tau_h \text{ for MMI V} = 10^{(5-1.7601)/1.992} = 42 \text{ cm/s}^2$$

One of the goals of the experiment is to identify what warning radius would give the best performance in the New Zealand context. A smaller warning radius decreases the likelihood of precautionary alerts and false alarms, but also decreases the warning times, and risks losing coverage if the sensor density is not high enough. The version of PLUM used in Japan had a warning radius of 30 km because the density of sensors is high. In Southern California, where the density of sensors is lower, they used a warning radius of 60 km (Saunders et al., 2022b). Increasing the warning radius increases coverage and warning times but will also increase the number of false alarms and violates the PLUM approximation of assuming undamped propagation. The analyses used a range of warning radii of 30 km, 40 km, 50 km, and 60 km to determine the best trade-off between coverage, warning times, and false alarms given the variable sensor density in New Zealand. The confirmation radius of 80 km is based on the density of sensors, as described in Section 4.2.5.

#### **4.2.7 EEW performance indicators**

The performance of the EEW system was evaluated based on key performance indicators that quantify alert accuracy and reliability at each grid location. Researchers have purposed different performance metrics to evaluate EEW performance. The framework initially proposed by Meier (2017) and extended in Meier et al. (2020) classifies alerts by comparing observed ground motion at a site with a predefined threshold. Under this framework, a True Positive (TP) occurs when an alert is issued and the observed shaking at the site exceeds the threshold. A False Positive (FP) occurs when an alert is issued, but shaking remains below the threshold. A False Negative (FN) is when no alert is issued, but the shaking exceeds the threshold, and a True Negative (TN) is when no alert is issued and the shaking remains below the threshold.

To capture a more detailed picture of EEW performance, additional classifications were introduced by Saunders et al. (2022a). These include Correct Alert (timely and late), Late Alert, Precautionary Alert and False Alert. A Correct Alert happens when you get a warning and then actually feel shaking at or above the threshold level—if the warning comes before the shaking it is a Timely Alert, and if it comes after, it is a Late Alert. A Missed Alert is when you do not get no warning but do feel strong shaking. A Precautionary Alert is when you get a warning

but the shaking stays below the level. A False Alert is when you get a warning, but no earthquake happens at all. In this research, we use similar categories to evaluate PLUM performance.

The PLUM algorithm’s performance was tested across various radii using the major performance indicators described below. The performance indicators for each single grid cell (g) over a 100-year period give a measure of the performance at different geographical locations across the region. The expected value of the performance indicator was computed by taking the mean over all the 100 periods of the value of the performance indicator in each period. The performance indicators were also summed or averaged over all the grid points for all the 100-year periods to give overall measures for the whole region.

*Table 4-4 Basic Performance Indicators for EEW performance at a single grid cell.*

<b>Indicator</b>	<b>Meaning</b>
$CTA_{g,p}$	Total number of Correct Timely Alerts received at a grid point. (alert received before strong shaking event ( $PGA \geq \tau_h$ ) at the grid point).
$LA_{g,p}$	Total number of Late Alerts received at a grid point (alert received after a strong shaking event at the grid point)
$MA_{g,p}$ :	Total number of Missed Alerts: no alert was received, but there was a strong shaking event at the grid point.
$MLA_{g,p}$	Missed or Late Alerts: Sum of $MA_{g,p}$ and $LA_{g,p}$
$PA_{g,p}$	Total number of precautionary alerts: alerts received at a grid point at which there was only a weak shaking event ( $PGA$ between $\tau_l$ and $\tau_h$ )
$FA_{g,p}$	Total number of false alerts: alerts received at a grid point when there was no noticeable shaking ( $PGA < \tau_l$ ) at the grid point

**Normalised measures:  $CTA_{norm}$ ,  $FA_{norm}$ , and  $PA_{norm}$**

Since the total number of shaking events and alerts at a grid point differ between grid cells and between the 100-year periods, the total number of CTA, PA, or FA counts at a grid point are difficult to compare, so normalized CTA values are more useful in evaluating the performances. As in chapter 3, the normalized CTA for a grid point in a particular 100-year period is the fraction of strong shaking events received by the given grid point for which a correct timely alert was received, as shown in equation 4.4.

$$CTA_{norm,g,p} = \frac{CTA_{g,p}}{(No. \text{ of shakings} \geq Th)_{g,p}} = \frac{CTA_{g,p}}{CTA_{g,p} + MLA_{g,p}} \quad (4.4)$$

where  $g$  is a grid location, and  $p$  is one of the 100 periods.

To obtain the expected value of  $CTA_{norm}$  at a grid cell from the results of all 100 sets of 100-year periods, the mean of the  $CTA_{norm}$  values is calculated over all the periods for which there was at least one strong shaking event at that grid cell, as shown in Eq 4.5. The same method can be applied for the MLA performance calculation.

$$\text{Expected } CTA_{norm,g} = \text{mean}_{p|SS_{g,p}>0} (CTA_{norm,g,p}) \quad (4.5)$$

where  $g$  is a grid location,  $p$  is one of the 100 periods,

$SS_{g,p}$  is the Number of strong shakings at grid location  $g$  in period  $p$ ,

and the mean is taken over all periods where  $SS_{g,p} \geq 0$ ,

In chapter 3, the false alerts were normalised by dividing by the number of non-strong-shaking events, but this led to a not very meaningful value, since it was affected by the number of grid cells where there was no shaking. A more useful approach for the false alerts is to normalise them relative to the total possible alerts at the same grid cell. This means that  $FA_{norm}$  represent the proportion of all possible alerts that were false . The same normalisation works for the precautionary alerts. Equations 4.6 and 4.7 define the value of  $FA_{norm}$  and  $PA_{norm}$  for a particular grid cell in a particular 100-year period.

$$FA_{norm,g,p} = \frac{FA_{g,p}}{(No. \text{ of all alerts})_{g,p}} = \frac{FA_{g,p}}{CTA_{g,p} + MLA_{g,p} + PA_{g,p} + FA_{g,p}} \quad (4.6)$$

$$PA_{norm,g,p} = \frac{PA_{g,p}}{(No. \text{ of all alerts})_{g,p}} = \frac{PA_{g,p}}{CTA_{g,p} + MLA_{g,p} + PA_{g,p} + FA_{g,p}} \quad (4.7)$$

Again, the expected values of  $FA_{norm}$  and  $PA_{norm}$  from all 100 sets of 100-year periods are obtained by calculating the mean of the  $FA_{norm,g,p}$  and  $PA_{norm,g,p}$  values over all periods for which there was at least one alert issued to the grid cell, as in eqs. 4.8 and 4.9

$$\text{Expected } FA_{norm,g} = \text{mean}_{p|AA_{g,p}>0} (FA_{norm,g,p}) \quad (4.8)$$

$$\text{Expected } PA_{norm,g} = \text{mean}_{p|AA_{g,p}>0} (PA_{norm,g,p}) \quad (4.9)$$

where  $g$  is a grid location,  $p$  is one of the 100 periods,

$AA_{g,p}$  is the number of all alerts at grid location  $g$  in period  $p$ . and

the mean is taken over all periods where  $AA_{g,p} \geq 0$

The measures above describe the performance at a single grid cell. To give a summary of the performance over the whole region, the measures are computed over the strong shaking events or alerts at all grid points and in all periods. The expected  $CTA_{norm}$  (eq 4.10) is therefore the probability of getting a correct timely alert when a strong shaking event occurs at some location in the region. The expected  $FA_{norm}$  or  $PA_{norm}$  (eqs 4.11 or 4.12) are the probability that an alert received at some location in the region is a false alert or a precautionary alert.

$$\text{Expected } CTA_{norm} = \frac{\sum_p(\sum_g(CTA_{g,p}))}{\sum_p(\sum_g(CTA_{g,p} + MLA_{g,p}))} \quad (4.10)$$

$$\text{Expected } FA_{norm} = \frac{\sum_p(\sum_g(FA_{g,p}))}{\sum_p(\sum_g(CTA_{g,p} + MLA_{g,p} + PA_{g,p} + FA_{g,p}))} \quad (4.11)$$

$$\text{Expected } PA_{norm} = \frac{\sum_p(\sum_g(PA_{g,p}))}{\sum_p(\sum_g(CTA_{g,p} + MLA_{g,p} + PA_{g,p} + FA_{g,p}))} \quad (4.12)$$

### **Warning time at a grid point and over region: WTCA and WTCTA**

An effective alert must give sufficient time for the recipients to act before the earthquake arrives. The warning time of an alert is the time between the alert time plus the communication and processing delay and the time of arrival of strong shaking. The alert time is the later of the time the alert message was sent and the time the confirmation message was sent. The warning time of an alert,  $a$ , is given by eq 4.13. The communication and processing delay was assumed to be 2 seconds (see Section 3.1).

$$WT_{a,g} = t_{shaking \text{ at } g \geq Th} - t_a - \text{comm\_delay} \quad (4.13)$$

where  $t_a$  is the time of the alert, and  $t_{shaking \text{ at } g \geq Th}$  is the time strong shaking is first detected at the grid point where the alert was received.

The warning time can only be computed for correct alerts (timely or late); for precautionary or false alerts, there is no strong shaking at the grid point, and the warning time is not defined. If the warning time is positive, then the alert was a Correct Timely Alert; if the warning time was negative, then the alert was a Late Alert.

We can characterise the warning times at a grid cell by the mean of the warning times of the alerts at that grid cell.

The mean warning time at a grid cell can be computed over all the correct alerts, including both timely and late alerts ( $WT_{CA}$ ). Alternatively, it can be computed just over the correct timely alerts ( $WT_{CTA}$ ). Although  $WT_{CTA}$  ignores the late alerts, it may be a more useful measure of how much warning is provided in the cases where there a timely warning to which users could respond. We therefore use both measures. Equations 4.14 and 4.15 give the two mean warning time measures for a single grid cell.

$$\text{Mean } WT_{CA_g} = \text{mean}_{a \in AlertsCA_g} (WT_{a,g}) \quad (4.14)$$

$$\text{Mean } WT_{CTA_g} = \text{mean}_{a \in AlertsCTA_g} (WT_{a,g}) \quad (4.15)$$

where  $g$  is a grid location,

$WT_{a,g}$  is the warning time of alert  $a$  received at grid location  $g$ ,

$AlertsCA_g$  is the set of all correct alerts at grid location  $g$  in all periods, and

$AlertsCTA_g$  is the set of all correct timely alerts at grid location  $g$  in all periods.

A summary over the whole region can be computed by finding the distribution of the warning times of all correct alerts (or all correct timely alerts) at all grid cells in all periods.

#### 4.2.8 Population-weighted measures

Section 3.4.3 of the previous chapter presented the need to weight the performance measures by population, based on the very uneven distribution of the population in the Canterbury region. However, that section multiplied the performance measures by the average population density in each cell, which led to performance measures that were hard to interpret, especially for the individual grid cells. A more appropriate method is to use the total population of each cell as a weighting. For the overall region measures, rather than taking the sum of the counts in each cell, it is more appropriate to weight the counts in a cell by the population of the cell.

The population for each grid cell was calculated using the population data of the region in 2020 (WorldPop, 2020). The population data is given for 30"x30" data cells. The population

of each grid cell was calculated from the sum of the populations of the data cells that fell within the grid cell.

Figure 4-10 illustrates the population distribution in the selected region, highlighting that Christchurch City has a significantly higher population compared to surrounding rural areas. Clearly, the alerts have more impact in Christchurch City than in the rural areas of the Northwest region, which is reflected in the population-weighted measures.




Figure 4-10 Population of the selected region in Canterbury. The colour of each square shows the population of the grid cell.

In assessing the performance of PLUM, we primarily focus on several key metrics: the fraction of strong shaking events for which a Correct Timely Alert is issued ( $CTA_{norm}$ ), and the fraction of alerts that are false alerts ( $FA_{norm}$ ) or precautionary alerts ( $PA_{norm}$ ). Therefore, the metrics are evaluated considering the population of each grid cell.

Equations 4.16 to 4.18 give the formulas for the population-weighted measures of the normalised CTA, FA and PA. These measures can be interpreted as the expected fraction of times a person in the region would receive a correct timely alert or a false alert or a

precautionary alert, rather than the fraction of times a grid cell would receive such an alert. Equations 4.19 and 4.20 give the formulas for the population-weighted measures of the warning time,  $WT_{CA}$  and  $WT_{CTA}$ , which can be interpreted as the average time that a randomly selected person in the region would need to wait between receiving an alert and experiencing strong shaking.

$$\text{Population-weighted } CTA_{\text{norm}} = \frac{\sum_p(\sum_g(CTA_{g,p} \times pop_g))}{\sum_p(\sum_g((CTA_{g,p} + MLA_{g,p}) \times pop_g))} \quad (4.16)$$

$$\text{Population-weighted } PA_{\text{norm}} = \frac{\sum_p(\sum_g(PA_{g,p} \times pop_g))}{\sum_p(\sum_g((CTA_{g,p} + MLA_{g,p} + PA_{g,p} + FA_{g,p}) \times pop_g))} \quad (4.17)$$

$$\text{Population-weighted } FA_{\text{norm}} = \frac{\sum_p(\sum_g(FA_{g,p} \times pop_g))}{\sum_p(\sum_g((CTA_{g,p} + MLA_{g,p} + PA_{g,p} + FA_{g,p}) \times pop_g))} \quad (4.18)$$

$$\text{Population-weighted } WT_{CA} = \frac{\sum_g(\sum_{a \in AlertsCA_g}(WT_{a,g} \times pop_g))}{\sum_g(|AlertsCA_g| \times pop_g)} \quad (4.19)$$

$$\text{Population-weighted } WT_{CTA_g} = \frac{\sum_g(\sum_{a \in AlertsCTA_g}(WT_{a,g} \times pop_g))}{\sum_g(|AlertsCTA_g| \times pop_g)} \quad (4.20)$$

Where  $pop_g$  is the population of grid cell  $g$ ,

$AlertsCA_g$  is the set of all correct alerts at grid cell  $g$  in all periods,

$AlertsCTA_g$  is the set of all correct timely alerts at grid location  $g$  in all periods, and

$WT_{a,g}$  is the warning time of alert  $a$  received at grid location  $g$ .


#### 4.2.9 Assessment of the EEW performance in simulated failure scenarios

Sections 4.2.7 to 4.2.8 describe the evaluation of the performance of the PLUM algorithm, assuming that all sensors and communication networks work perfectly. However, during a major earthquake event, there is a possibility that the sensors, networks, or both could fail, which would affect the EEW system. Therefore, it is also necessary to evaluate the performance under simulated failure scenarios.

There are several possible types of failure scenarios, ranging from minor to catastrophic. For example, one sensor could fail, or a number of sensors could fail, a network link could disappear, a network provider could fail, or the main central system could crash. We will first consider the major but not catastrophic scenario where one of the network service providers has a breakdown and then consider a worse scenario where the sensors in the band around

Christchurch all failed. In each case, we measure the performance metrics as in the previous analyses and compare them to the fully working case.


GeoNet currently utilises three service providers for data communication from the sensors (R. Guest, personal communication, 2021; GeoNet, 2023) For the first experiment, we make the simplifying assumption that one-third of the stations are connected to each service provider and these three service providers are randomly distributed among the sensors. A breakdown of one service provider disconnecting the stations connected to that service provider from the network can therefore be represented by removing one-third of the stations randomly as illustrated in Figure 4-11.



*Figure 4-11 Failure Scenario A: removing one-third of the stations randomly to represent a breakdown of one service provider. Triangles represent sensors. Red triangles are the removed sensors, Black triangles are the remaining sensors. The red square represents the Canterbury region selected to evaluate performance. The black dotted circle represents the 250 km radius circle containing all the sensors that are relevant to the selected Canterbury region.*

In the second experiment, we removed stations in a band around Christchurch between 30 km and 60 km from Christchurch City, as illustrated in Figure 4-12. For this scenario, the centre

of the most highly populated grid cell of the Canterbury region is considered as the centre of the city. We analysed the performance as in the first experiment.



*Figure 4-12 Failure Scenario B: removing stations in a band between 30 km and 60 km from the centre of Christchurch City. Blue circles show the 30 km and 60 km limits. Triangles represent sensors. Red triangles are the removed sensors, Black triangles are the remaining sensors. The red rectangle represents Canterbury region selected to evaluate performance. The Black dotted circle represents the 250 km radius circle containing all the sensors that are relevant to the selected Canterbury region.*

### 4.3 Results for the Canterbury region

The PLUM algorithm was tested using the generated synthetic seismograms of all 945 earthquakes across 182 grid locations and 131 station locations. The PLUM algorithm described in Section 4.2.6 was applied to each earthquake, and the results were aggregated to present the overall result for each 100-year period. Finally, the results were averaged over

all 100 periods to obtain the expected measures of the performance of an EEW system utilising PLUM over the next century if it were built on the current set of GeoNet sensors.

First, the results are presented without considering the population, and then presented taking the population into consideration.

#### 4.3.1 Results for the Canterbury region by location

The algorithm was evaluated on each 100-year period, applying it to each selected earthquake of that period, and for each different warning radius.

Figure 4-13 shows results for earthquakes of one particular 100-year period (#69) for the given parameters (warning radius =50 km,  $\tau_h=42$  cm/s<sup>2</sup>,  $\tau_l=13$  cm/s<sup>2</sup>). The eight earthquakes of the period were highlighted as red circles in Figure 4-7. The size of the coloured square in each grid cell of Figure 4-13 represents the number of strong shaking events at that grid location resulting from the earthquakes. For this particular period, it is clear that the earthquakes caused the most shaking events in the north-west corner of the region, some shaking events in the south-west quarter of the region but few shaking events in the eastern half of the region. Note that the site class at the grid point affects the shaking intensity, so there may be shaking exceeding threshold at a grid points of site E, even when shaking does not reach the threshold at neighbouring grid points of a different site class.

The colours of the squares in Figure 4-13 indicate the fraction of correct timely alerts ( $CTA_{norm,g}$ ) for those strong shaking events. Dark blue means 100% CTA (the given grid point received correct timely alerts for 100% of the shaking events). Dark red means 0% CTA (in other words 100% MLA). 50% CTA is represented by the light colour between pink and blue.

The figure shows that for these 8 earthquakes, PLUM generated mostly correct timely alerts except in the middle of the western border of the region. However, many grid cells had no strong shaking events over the whole period, including much of the Christchurch region, and many had only one event. While looking at a single period is helpful, there are too few earthquakes to make any reliable conclusions. It is essential to consider the performance averaged over all the periods to obtain meaningful results.




Figure 4-13 Fraction of shaking events with Correct Timely Alerts at each grid point for a selected (period #69) 100-year period, PLUM radius =50 km,  $\tau_h=42$  cm/s<sup>2</sup>,  $\tau_l=13$  cm/s<sup>2</sup>. The colour of each square shows the value of  $CTA_{norm,g}$  where dark blue is 100% CTA and dark red is 100% MLA. The size of each square represents the number of shaking events above the  $\tau_h$  threshold received by the grid point (smallest square=1, largest square=3). Insert shows the positions of the sensors in the region (see Figure 4-3 for all sensors). See Figure 4-7 for all earthquake sources.

Figure 4-14 illustrates the expected  $CTA_{norm,g}$  performance (see eq. 4.5), based on all 100 periods, using a warning radius of 50 km. Most of the grid cells of the selected region experienced some shaking above the threshold in at least some of the periods. Comparing the average results in Figure 4-14 to the results of a single 100-year period in Figure 4-13 indicates that there is a possibility of getting strong shaking (above the threshold) for almost all locations across the region. However, most of the shaking occurs in the Northwest region because most of earthquakes are closer to that area. There is less shaking in the Southeast region, indicated by the smaller size of the squares, or no shaking at all indicated by blank grid cells.

Consideration of Figure 4-14 suggests that, at a 50 km warning radius, Christchurch City almost always receives correct timely alerts: the grid cell covering the centre of the city (at -

43.5° to -43.6°, 172.6° to 172.7°) achieved 100% CTA. In contrast, the west region of the area experiences more missed and late alerts.

However, the population of the rural areas in the west is much less than the Christchurch City area, and Section 4.3.2 will discuss how to consider population density to evaluate the performance more usefully.

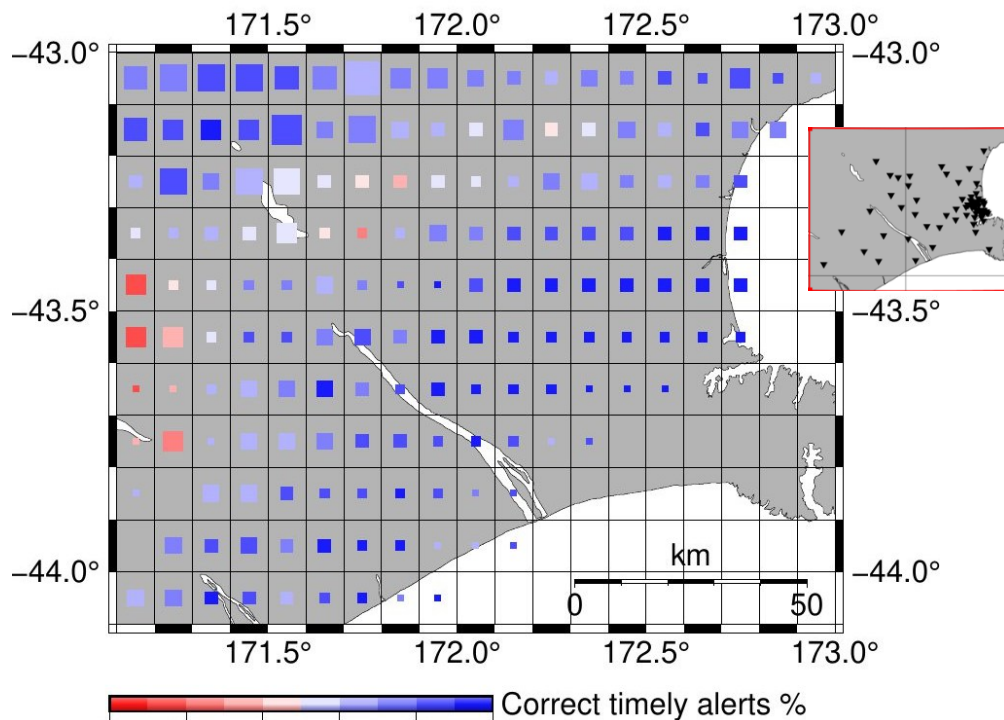


Figure 4-14 Expected values of  $CTA_{norm,g}$  for all 100 random 100-year periods, PLUM radius = 50 km,  $T_h=0.42 \text{ m/s}^2$ ,  $T_l=0.13 \text{ m/s}^2$ . Size of squares shows expected number of shaking events above the  $\tau_h$  (smallest square=0.1, largest square=2). Colour of squares shows  $CTA_{norm,g}$ , blue = 100% CTA; red = 0% CTA (100% MLA). Insert shows the positions of the sensors (see Figure 4-3 for all sensors). See Figure 4-7 for all earthquake sources.

Figure 4-15 shows the values of the expected  $FA_{norm,g}$  (fraction of alerts that are false alerts) in each grid cell. The size of the squares represents the expected number of all alerts (CTA+MLA+PA+FA) received by each grid cell.

All the grid cells of the region get less than 19 % false alerts, with the majority of grid cells reporting a false alerts percentage less than 1 % (101 grid locations out of 182). Only 14 grid locations get more than 10% False Alerts.

The squares around Christchurch City are larger, showing the total number of alerts are higher around the city. This is due to the high station density. Most of those locations also get a smaller number of false alerts.

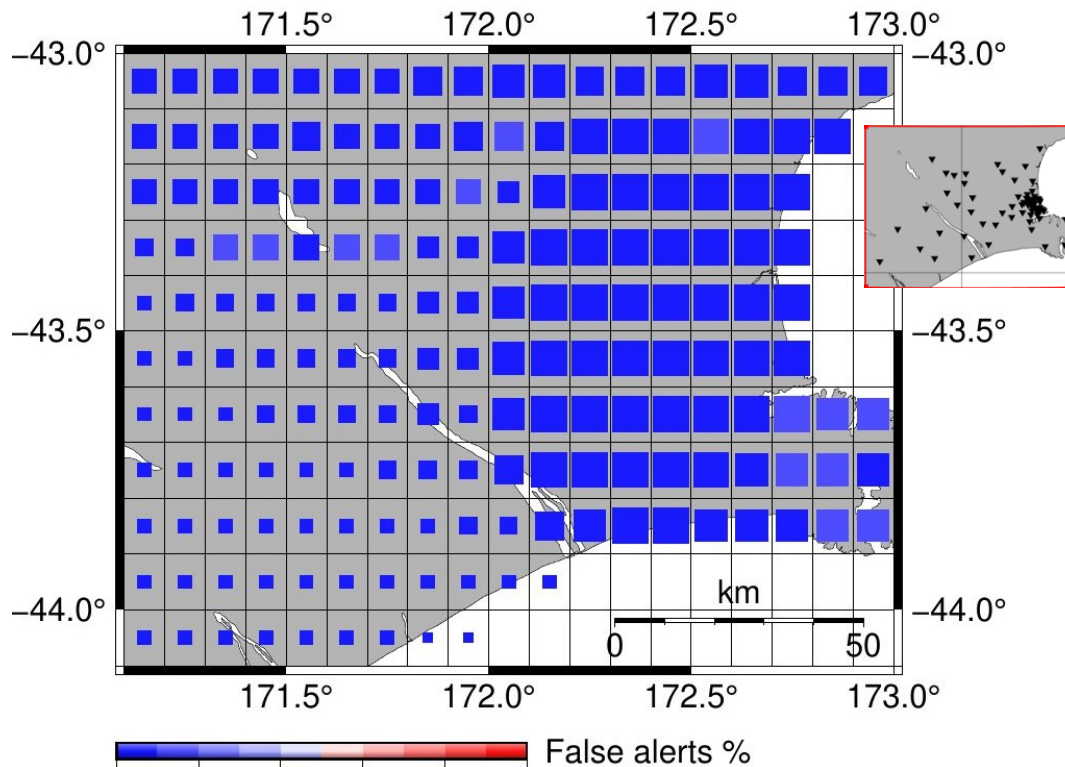


Figure 4-15 Expected values of False Alerts vs All Alerts (expected  $FA_{norm,g}$ ). PLUM radius =50 km,  $T_h=0.42 \text{ m/s}^2$ ,  $T_l=0.13 \text{ m/s}^2$ . Size of squares shows expected number of all alerts received by the grid cell (smallest square= 0.6, largest square = 6.2). Colour of squares shows expected  $FA_{norm,g}$ , blue = 0% FA; red = 100% FA. Insert shows the positions of the sensors in the region (see Figure 4-3 for all sensors). See Figure 4-7 for all earthquake sources.

Figure 4-16 illustrates expected  $PA_{norm,g}$ , the number of Precautionary Alerts as a percentage of the number of all alerts received by that grid location. As the count of all alerts is higher around the Christchurch City, that area receives more precautionary alerts. The figure shows that most alerts will be precautionary. It is only in some grid cells in the west of the region that less than 50% of the alerts are precautionary alerts.


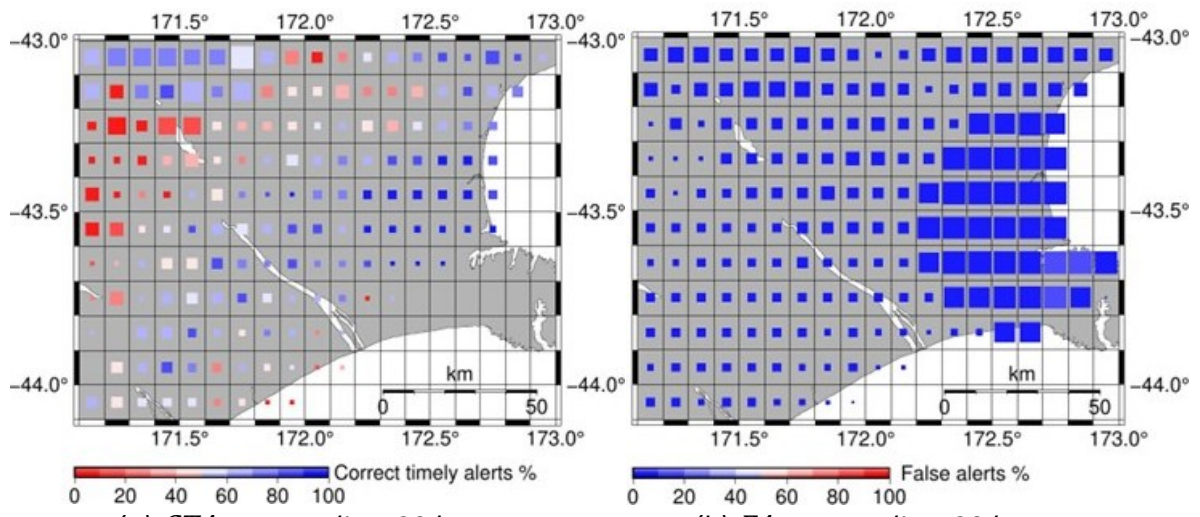


Figure 4-16 Expected value of Precautionary Alerts Vs All Alerts (expected  $PA_{norm,g}$ ). PLUM radius = 50 km,  $T_h=0.42 \text{ m/s}^2$ ,  $T_l=0.13 \text{ m/s}^2$ . Size of squares show expected number of all alerts received by the grid cell (small square=0.6, largest square=6.2). Colour of squares shows expected  $PA_{norm,g}$ , blue = 0% PA; red = 100% PA. Insert shows the positions of the sensors in the region (see Figure 4-3 for all sensors). See Figure 4-7 for all earthquake sources.

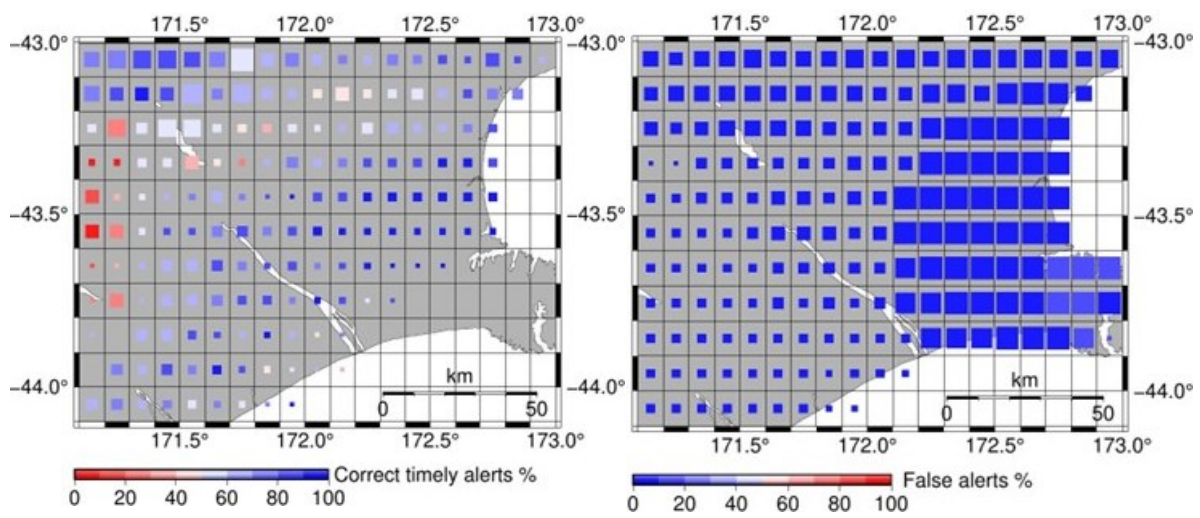
#### Grid cell measures: $CTA_{norm}$ , $FA_{norm}$ and $PA_{norm}$

Figure 4-17 illustrates the relationship between the warning radius and the two key performance measures,  $CTA_{norm}$  and  $FA_{norm}$ . The figure shows that the higher radius will be best with respect to the CTA performance and worst for False Alerts. However, further analysis is required to choose the optimum warning radius considering other alerts such as Precautionary Alerts and population.



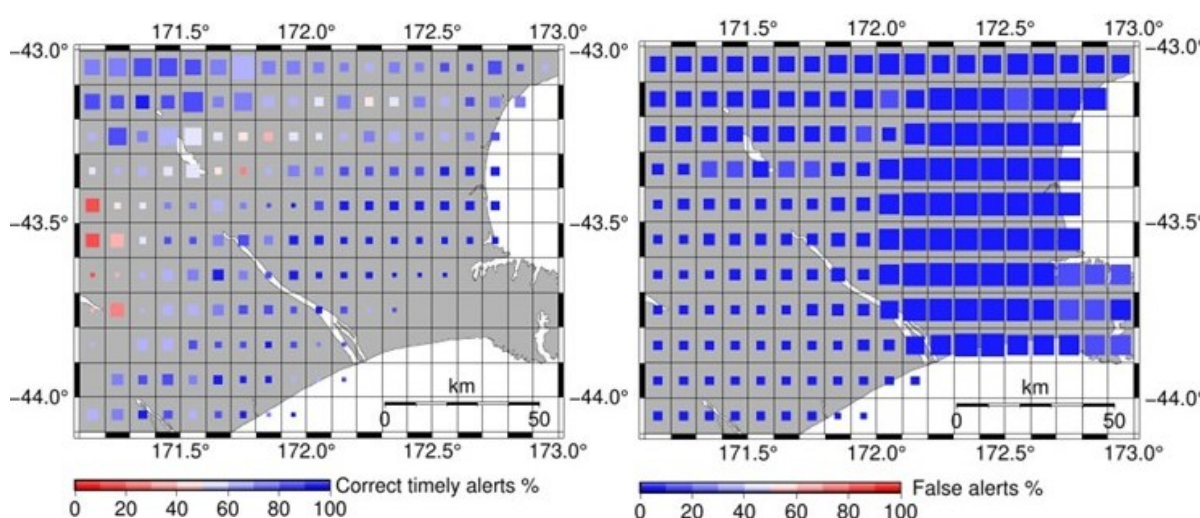
(a)  $CTA_{norm,g}$ : radius=30 km

(b)  $FA_{norm,g}$ : radius=30 km



(c)  $CTA_{norm,g}$ : radius=40 km

(d)  $FA_{norm,g}$ : radius=40 km



(e)  $CTA_{norm,g}$ : radius=50 km

(f)  $FA_{norm,g}$ : radius=50 km

(Figure continued on next page)

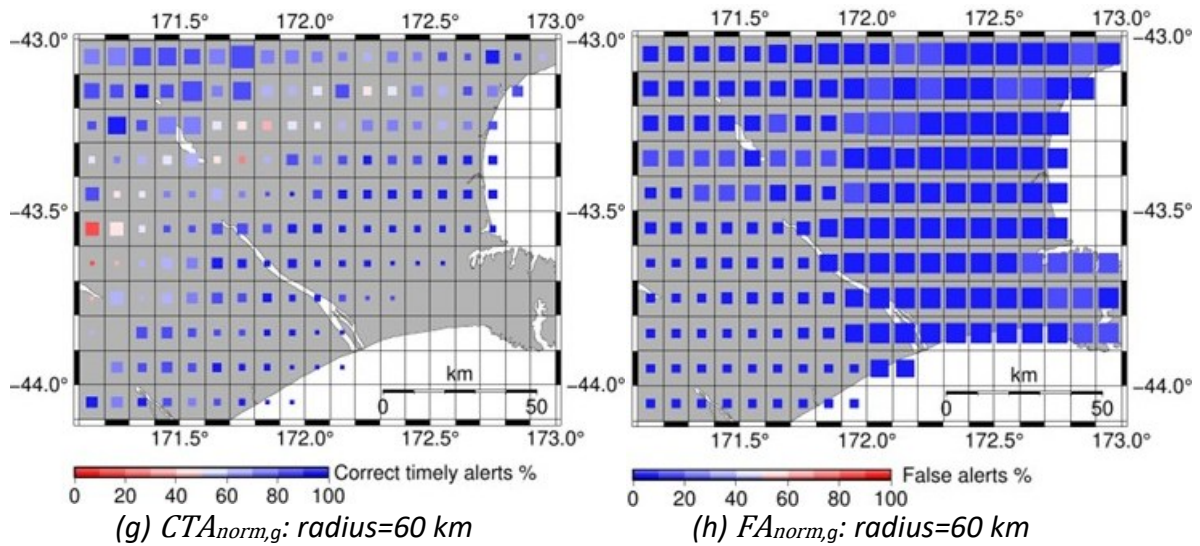


Figure 4-17 Trade-off between expected  $CTA_{norm,g}$  and expected  $FA_{norm,g}$  for different warning radii: 30 km, 40 km, 50 km, and 60 km.  $T_h=0.42 \text{ m/s}^2$ ,  $T_l=0.13 \text{ m/s}$ . Colour of squares as in Figure 4-14 (for a,c,e,g) and Figure 4-15 (for b,d,f,h). Square sizes as in Figure 4-14 for a,c,e,g and similar to Figure 4-15 (i.e., show the relative values and are similar but vary slightly for each figure) for b,d,f,h.

The number of precautionary alerts also increases with the warning radius as shown in Figure 4-18. However, the precautionary alert percentages around the Christchurch City remain similar at all four warning radii and therefore, the analysis that takes the population into consideration will be a better option for selecting an appropriate warning radius for the Canterbury region.

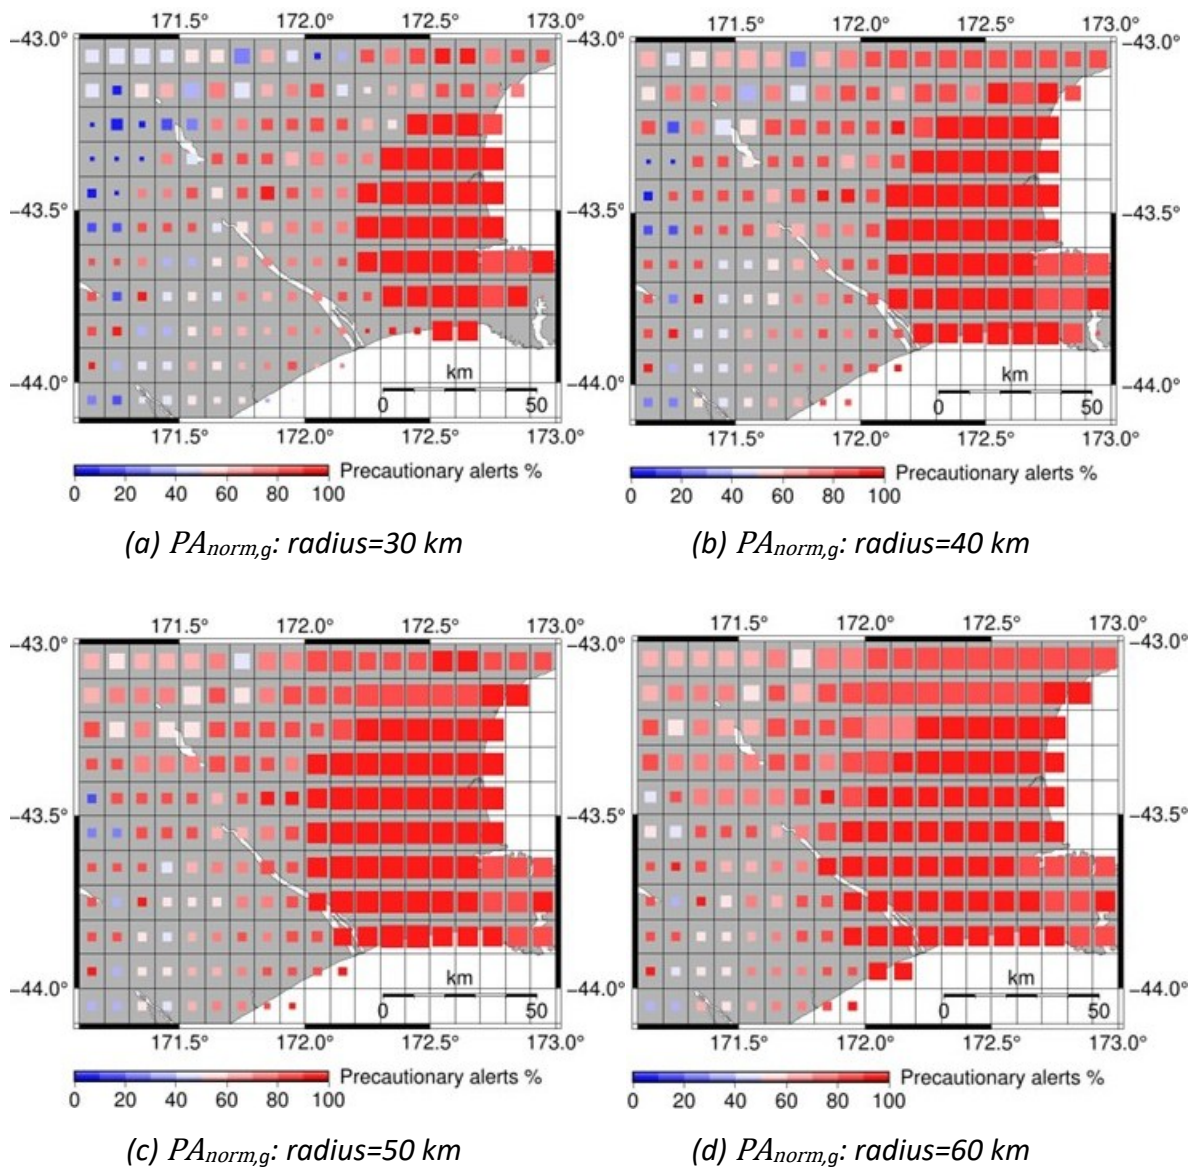


Figure 4-18 Values of Precautionary Alerts Vs All Alerts ( $PA_{norm,g}$ ), for different warning radii. PLUM radius 30,40,50,60 km,  $T_h=0.42$  m/s<sup>2</sup>,  $T_l=0.13$  m/s. Size of squares shows expected number of all alerts received by the grid cell. Square sizes show the relative values and are similar to Figure 4-16 but vary slightly for each figure. Colour of squares shows expected  $PA_{norm,g}$ , blue = 0% PA; red = 100% PA.

### Overall region measures

The analysis above considered the PLUM outcomes at each grid cell separately. It is also useful to consider the performance of PLUM over the region as a whole. Table 4-5 shows the expected total number of each alert outcome over 100 years at the different warning radii (30 km, 40 km, 50 km and 60 km). The values were obtained by determining the mean over all the 100 periods of the total number of each alert outcome (see the performance indicators

in Table 4-2) across all grid cells. In general, a higher warning radius generates more correct timely alerts, fewer missed/late alerts, but many more precautionary alerts and false alerts.

*Table 4-5 Expected number of each PLUM outcome across whole region for different warning radii.*

Warning radius	CTA (n)	CTA (%)	LA (n)	LA (%)	MA (n)	MA (%)	MLA (n)	MLA (%)	PA (n)	PA (%)	FA (n)	FA (%)	Total
30 km	37	9%	17	4%	13	3%	31	8%	291	74%	4	1%	393
40 km	46	9%	16	3%	6	1%	21	4%	434	81%	10	2%	533
50 km	50	8%	15	2%	2	0%	17	3%	540	84%	20	3%	644
60 km	53	7%	13	2%	1	0%	14	2%	642	85%	36	5%	759

It is clear from Table 4-5 that there is a significantly larger number of precautionary alerts (PA) than other types of alerts. For a 30 km warning radius, the number of PA is approximately eight times greater than the number of CTA. For 60 km it is around 12 times.

The number of false alerts (FA) is reasonable for the 30 km as it is only 11% of the CTA. However, when the radius increases, FA increases to 22%, 40%, and 68% as much as the CTA respectively. The largest improvement in CTA is when increasing the warning radius from 30 km to 40 km; changing to 50 km or 60 km has much less effect on the CTA, though the FA keeps rising with each step. This suggests that the 40 km radius may be the best trade-off.

As the warning radius increases, the number of late alerts (LA) and missing alerts (MA) decreases. For 30 km, MLA is equal to 84% of CTA and for the 60 km it is only 26%. Most of the new CTA generated when the warning radius is increased from 30 km to 40 km come from previously missing alerts. Extending the warning radius allows stations to provide alerts to grid points that received no alert at a smaller radius. The reduction in late alerts as the radius increases is much smaller.

Table 4-6 Table 4-6 shows the expected normalised performance measures for the whole region. The relevant equations are shown in Section 4.2.7 as equation 4.10, 4.11 and 4.12. The expected  $CTA_{norm}$  defined by equation 4.10, represents the probability of receiving a correct timely alert when a strong shaking event occurs at some location in the region. The expected  $CTA_{norm}$  is around 55% for the 30 km case and shows significant improvement when

the radius increases from 30 km to 40 km. Expected  $FA_{norm}$  represents the probability that an alert received at some location is a false alert. Expected  $FA_{norm}$  increases by 0.8% for 30 km to 40 km and then by 1.4% for 40 km to 50 km. Expected  $PA_{norm}$  represents the probability that an alert received at some location is a precautionary alert. It is significant that the expected  $PA_{norm}$  is high for all cases. According to these results, a warning radius between 40 km and 50 km would be a good parameter choice, considering the trade-off between  $CTA_{norm}$  and  $FA_{norm}$ . However, if the 4.9%  $FA_{norm}$  and 86.1%  $PA_{norm}$  is acceptable, then the 60 km warning radius would be a better choice because it gives a better CTA to MLA ratio.

*Table 4-6 Expected  $CTA_{norm}$ ,  $FA_{norm}$  and  $PA_{norm}$  across whole region for different warning radii. Note, the numbers in each row should not add to 100%;  $CTA_{norm}$  is a fraction of strong-shaking events, and  $FA_{norm}$  and  $PA_{norm}$  are fractions of alerts received.*

Warning radius	Expected $CTA_{norm}$	Expected $FA_{norm}$	Expected $PA_{norm}$
30 km	54.7%	1.1%	80.3%
40 km	68.1%	1.9%	84.9%
50 km	74.4%	3.3%	86.0%
60 km	79.1%	4.9%	86.1%

### Warning time measures

To be effective, an alert must not only be correct and timely, but it must give sufficient warning time for the recipients to take action before the earthquake arrives. Figure 4-19 shows the expected mean warning time for each grid cell, at the four different warning radii. Even at the 30 km warning radius, the majority of grid cells have positive mean warning times, indicating more timely alerts than late alerts, especially in the South-East half of the region. As the warning radius increases, the fraction of timely alerts increases and at 60 km warning radius, only a small band of grid cells still have a negative mean warning time. The maximum and minimum mean warning time values are displayed for each warning radius in captions (a)–(d) of Figure 4-19.

We also measured the warning time across the whole region, computing not only the mean warning time of all the correct alerts, but also the distribution of the warning times. Figure 4-20 and Figure 4-21 show the distribution of warning times.

Figure 4-20 illustrates the expected warning time distribution for the region considering the correct timely alerts only. Each box plot shows the median (red line), along with the 25th and 75th percentiles, and includes outliers to indicate the range of variation in warning times. The analysis already accounts for a 2-second communication delay. The median warning time for a 30 km radius is 5.3 seconds, increasing to 7.0 seconds, 8.3 seconds, and 9.3 seconds as the PLUM warning radius expands to 40 km, 50 km, and 60 km, respectively.

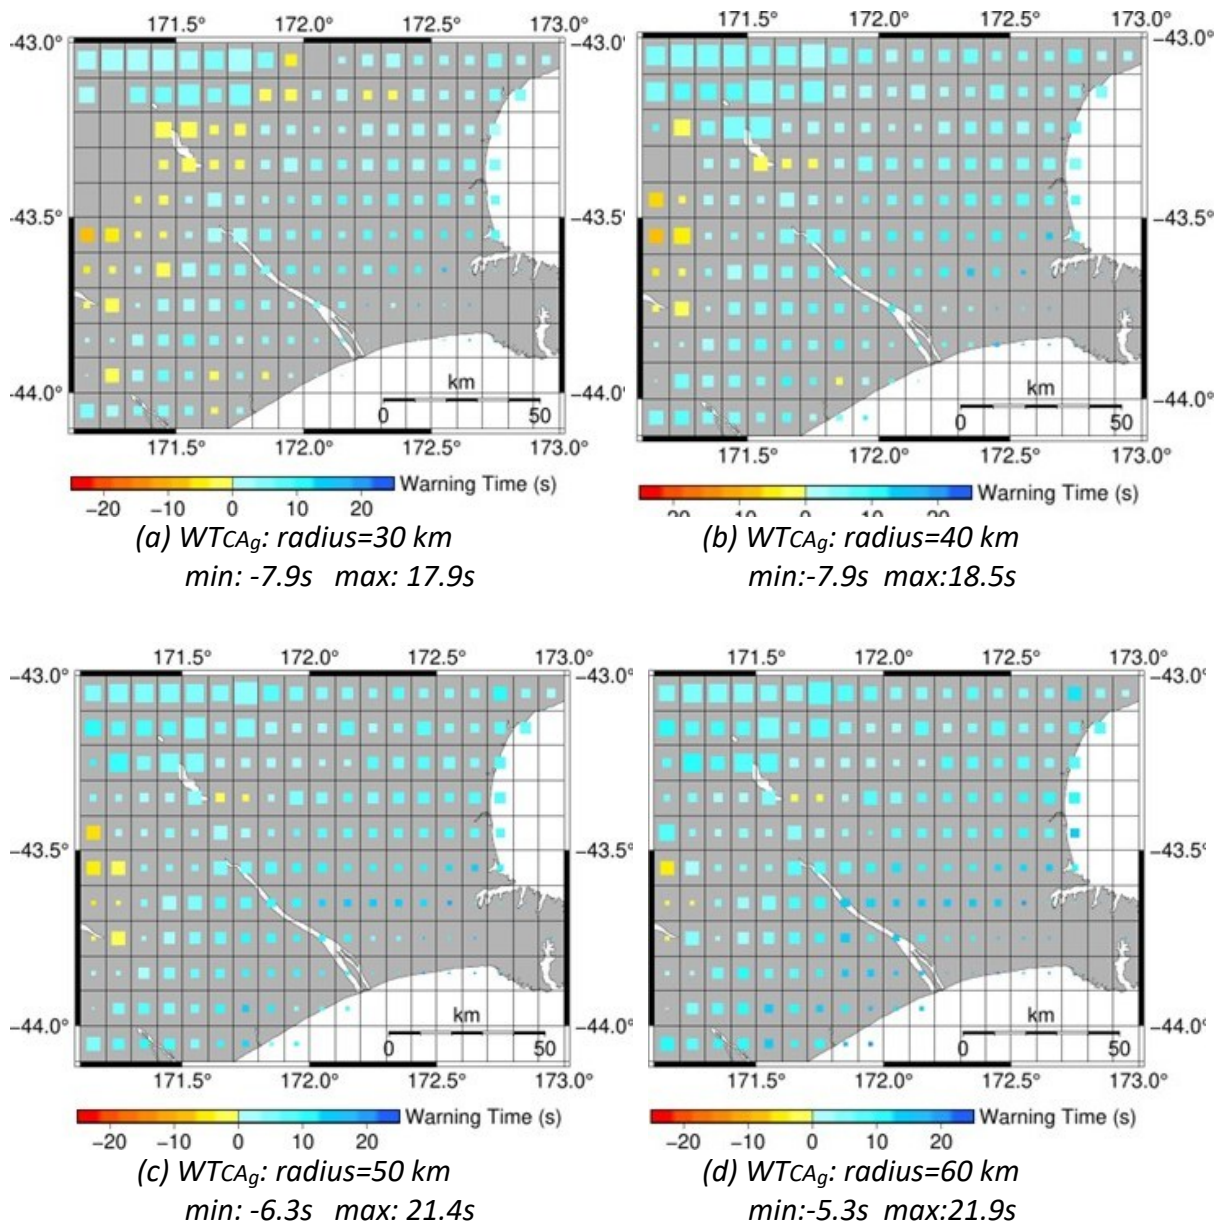


Figure 4-19 Expected mean warning times for correct alerts, for different warning radii (30 km, 40 km, 50 km, 60 km). Size of square represents the expected number of correct alerts in the cell (small square: 0.1, largest square: 1.6 (30 km), 1.6 (40 km), 1.8 (50 km), 2.0 (60 km)). Colour of square represents the expected mean warning time: yellow/orange for a negative mean warning time (late); blue for a positive mean warning time (timely). The sub-captions for each case show minimum and maximum mean warning times across all the cells.

Figure 4-20 Figure 4-20 does not include the negative warning times of the late alerts, whereas Figure 4-21 Figure 4-21 illustrates the distribution of expected warning times of all correct alert across the region including both timely and late alerts. The median warning time for a 30 km radius is 2.7 seconds, increasing to 4.5 seconds, 6.0 seconds, and 7.4 seconds as the PLUM warning radius expands to 40 km, 50 km, and 60 km, respectively.




Figure 4-20 Distribution of *WTCTA*: box plots of warning times of all correct **timely** alerts across region. Warning radii 30 km, 40 km, 50 km, 60 km. The box plots show the median (red line), 25<sup>th</sup> and 75<sup>th</sup> percentiles, and the outliers. The dotted lines show the 25<sup>th</sup> percentile, median, and 75<sup>th</sup> percentile of the 50 km warning radius for comparison with other warning radii.

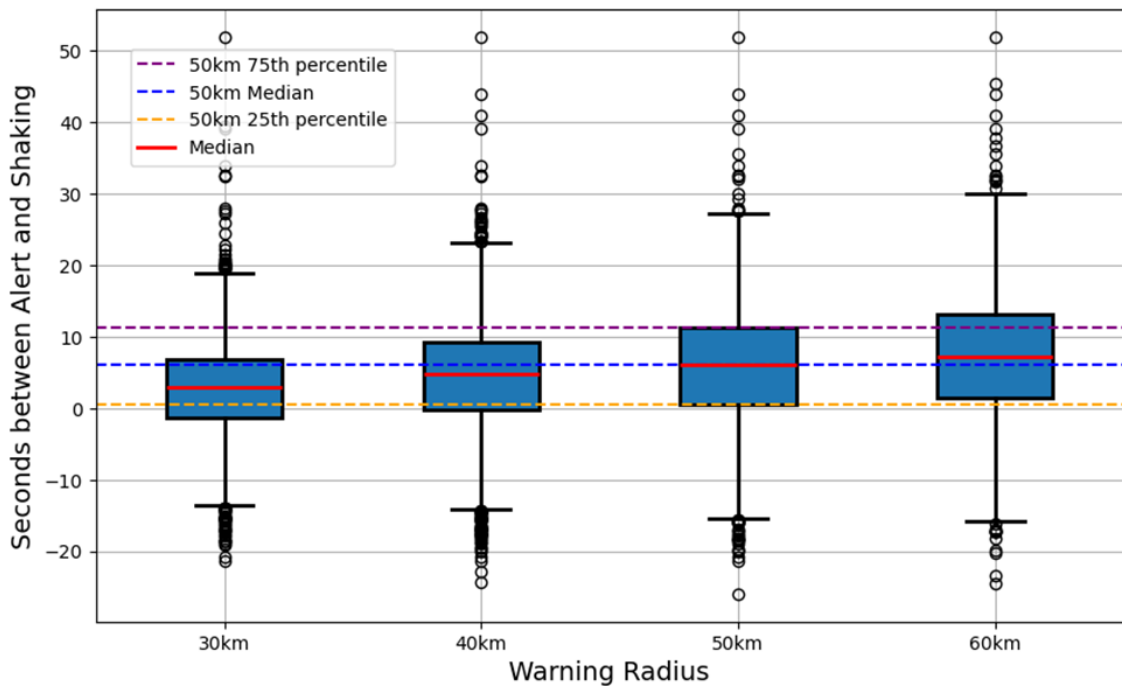


Figure 4-21 Distribution of *WTCA*: box plots of warning times of all correct alerts across region. Warning radii 30 km, 40 km, 50 km, 60 km. The box plots show the median (red line), 25<sup>th</sup> and 75<sup>th</sup> percentiles, and the outliers. The dotted lines show the 25<sup>th</sup> percentile, median, and 75<sup>th</sup> percentile of the 50 km warning radius for comparison with other warning radii.

The mean and standard deviation of the warning time distributions are shown in Table 4-7. To provide sufficient warning time for the region, a PLUM radius of 50 km or 60 km may be required.

*Table 4-7 Mean ± standard deviation of distributions of WTCTA (warning time for **timely** alerts only) and WTCA (warning time for all correct alerts) for different warning radii.*

Warning Radius	WTCTA (s)	WTCA (s)
30 km	6.0 ± 4.3	2.7 ± 6.3
40 km	7.7 ± 5.1	4.5 ± 7.2
50 km	9.0 ± 5.8	6.0 ± 7.8
60 km	10.2 ± 6.5	7.4 ± 8.2

These results suggest that PLUM EEW will perform well for the Christchurch City area but there are missed or late alerts in North-West rural areas of the selected region. However, to get performance scores that correspond better to the risk of possible injury and damage, it is necessary to take the population into account, which will be discussed in the next section.

### 4.3.2 Results for the Canterbury region by population

Section 4.3.1 described the performance of PLUM for the Canterbury region. However, these results did not account for the population density within the grid cells. Performance evaluation was solely based on the number of alerts and shaking events affecting relevant grid cells. Therefore, this section discusses the performance indicators, taking into account the population within the grid cells.

#### Population-weighted grid cell measures

Figure 4-22 shows the expected  $CTA_{norm,g}$  for each grid cell taking into account the population, with a warning radius of 50 km. Figure 4-23 shows the same results for 30, 40, 60 km radii. The colours of the squares show the same ratio of CTA to strong shaking events as in Figure 4-14. However, the areas of the squares in Figure 4-22 are weighted by the population, as well as the number of strong shaking events (greater than  $\tau_h$ ) experienced at that grid location. The area of the squares therefore represents the expected number of “person-events” for shaking events in that grid cell. Note that to ensure

that the colour of squares in low population cells is visible, squares that would have an area less than 1% of the largest square are shown as small fixed sized circles.

For the 50 km scenario of Figure 4-22, the grid cell with the largest square (close to Christchurch City) has a population of 132,399 and experienced 19 shaking events over the 100 sets of 100-year periods (an average of 0.2 events per 100 years). Therefore, the expected number of person-events at that grid cell is 26,480. The second largest square is 18,104 expected person-events (population 90522 x 0.2 events). For Figure 4-22, the small circles represent fewer than 250 person-events so the grid cells with squares represent the significant grid cells for EEW as the population highly weights them. The grid cells with the circles are much less critical due to the low population.

The expected  $CTA_{norm,g}$  performance of the significant high-populated locations around Christchurch City is similar for all four warning radii, as illustrated in Figure 4-14 and Figure 4-17.

The grid cells with the highest and second highest population achieved a 100% CTA rate, whereas all the grid cells with a low CTA rate (white, pink, or red) had low populations and fewer than 250 person-events. A rough measure of population-based performance is to examine the CTA rates just over the significant grid cells, indicated by squares, which gives a minimum CTA rate of 67% and a median of 93% and a mean of 90%. This indicates that PLUM would work well in the areas of high population density. This can be compared to the overall median and mean CTA rate across all grid cells of 79% and 77% respectively when population is ignored.

Note that this analysis assumes that the population distribution will be similar for the next 100 years.

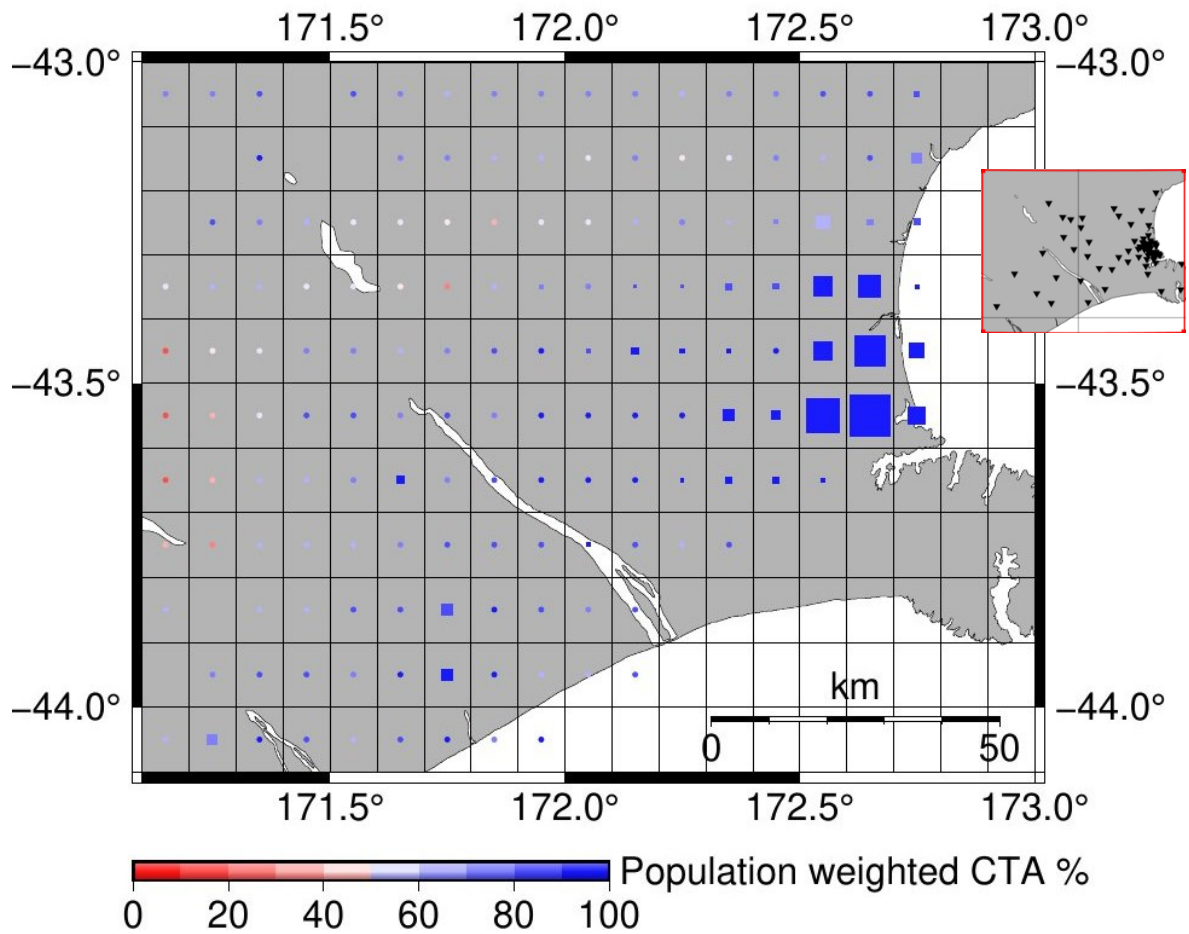


Figure 4-22 Population-weighted  $CTA_{norm,g}$  PLUM warning radius=50 km,  $\tau_h=0.42 \text{ m/s}^2$ ,  $\tau_l=0.13 \text{ m/s}^2$ . The area of each square represents the number of person-shaking events—the product of the population and the expected number of significant shaking events ( $pga \geq \tau_h$ ) experienced at that grid location (smallest square=250, largest square=26480). Color of each square represents the value of  $CTA_{norm,g}$ . To ensure that the colour of squares in low population cells is visible, squares that should have area less than 1% of the largest square are shown as small fixed-size circles. (For this case, circles denote those cells with fewer than 250 person-events. Insert shows the positions of the sensors in the region (see Figure 4-3 for all sensors).

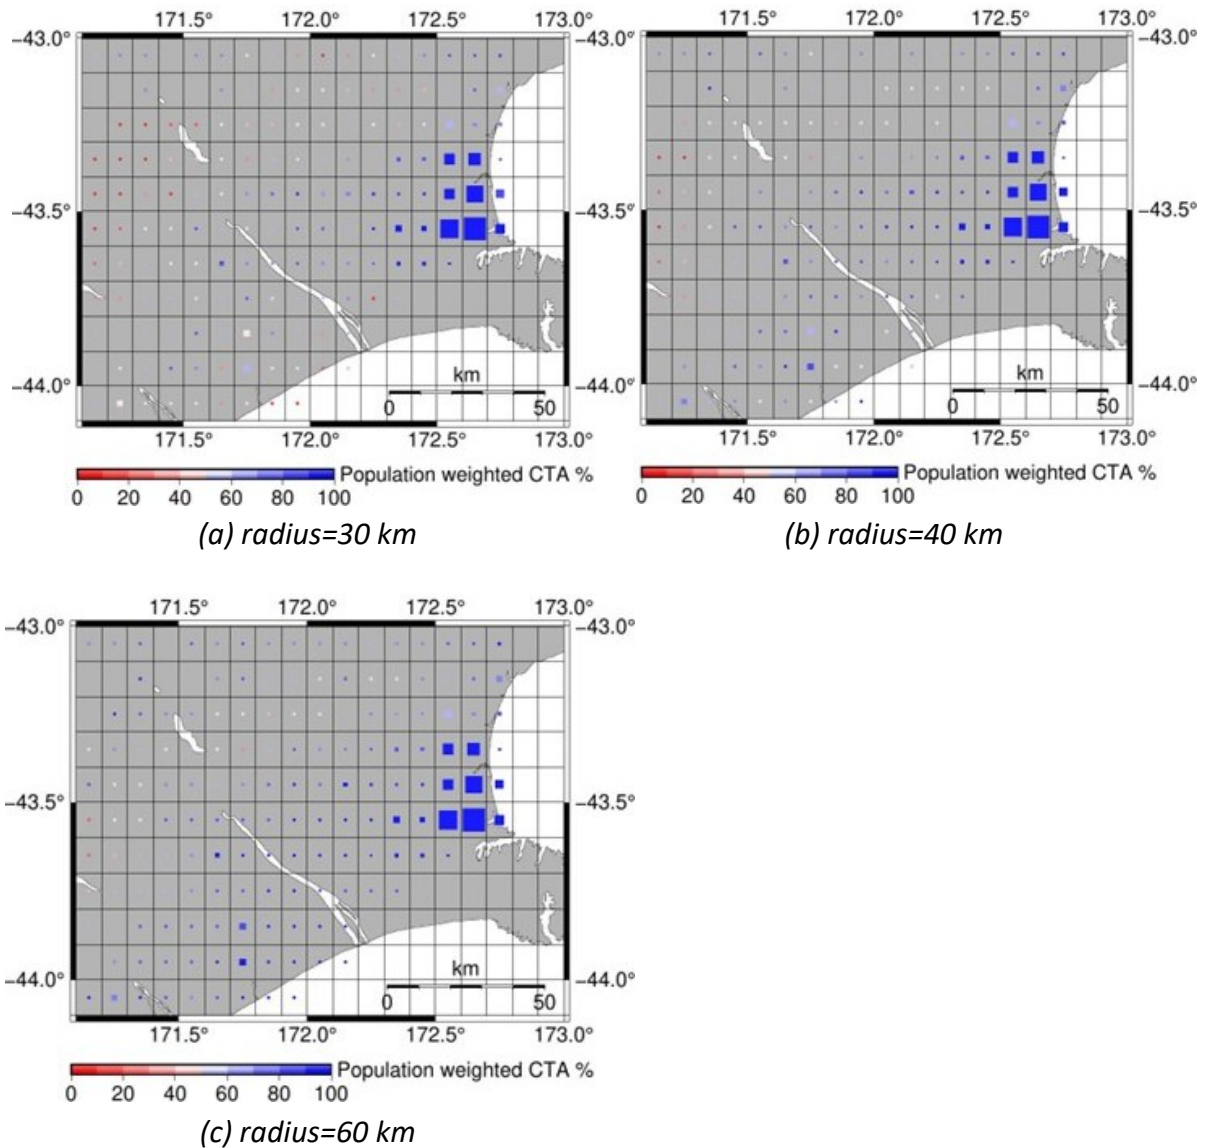


Figure 4-23 Population-weighted  $CTA_{norm,g}$ , other PLUM warning radii (a) 30 km, (b) 40 km, (c) 60 km. Other parameters are as in Figure 4-22.

Figure 4-24 Figure 4-24 shows the expected  $FA_{norm,g}$  for each grid cell taking into account the population, with a warning radius of 50 km. The colours of the squares show the same ratio of false alerts to all alerts as Figure 4-15. Since  $FA_{norm,g}$  is expressed as a fraction of the total number of alerts, the area of each square in Figure 4-15 is weighted by the population multiplied by the expected number of alerts received at that grid location. The area of the squares therefore represents the expected number of “person-alerts” received in that grid cell.

Figure 4-25 Figure 4-25 shows the expected  $FA_{norm,g}$  for 30, 40, 60 km warning radii and the results are similar to Figure 4-24 for the significant high populated grid cells.

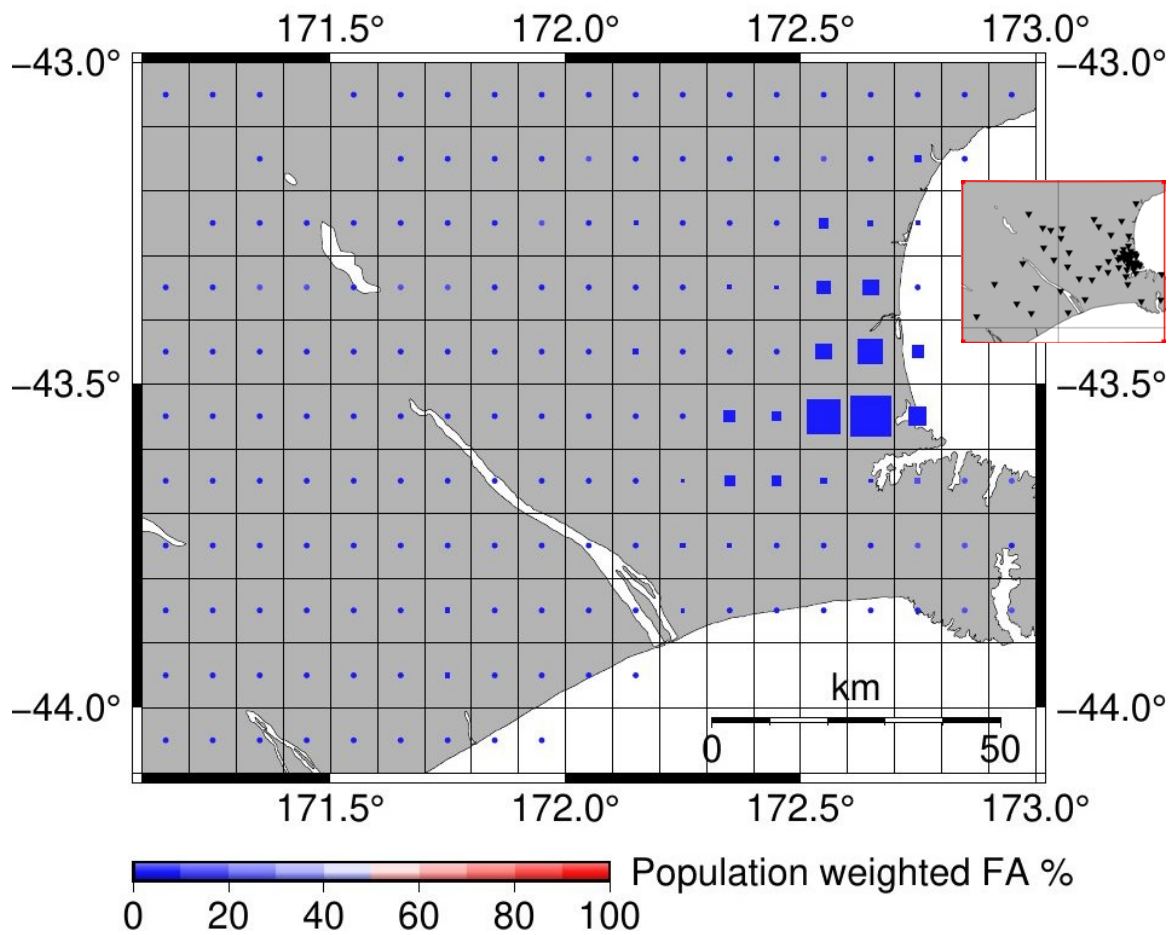


Figure 4-24 Population-weighted  $FA_{norm,g}$  PLUM radius =50 km: False Alerts vs. All Alerts for all 100 random 100-year periods (total 10,000 years),  $Th=0.42 \text{ m/s}^2$ ,  $Tl=0.13 \text{ m/s}^2$ . Colour represents the  $FA_{norm,g}$ : blue = 0% false alerts, red = 100% false alerts. The area of each square represents the number of person-alerts—the product of the population and the expected number of alerts received at that grid location (largest square=1853586, smallest square=6600). Squares that would have an area less than 1% of the largest square (ie, less than 6,600 person-alerts in this case) are shown as small fixed-size circles. The insert shows the positions of the sensors (see Figure 4-3 for all sensors) in the region.

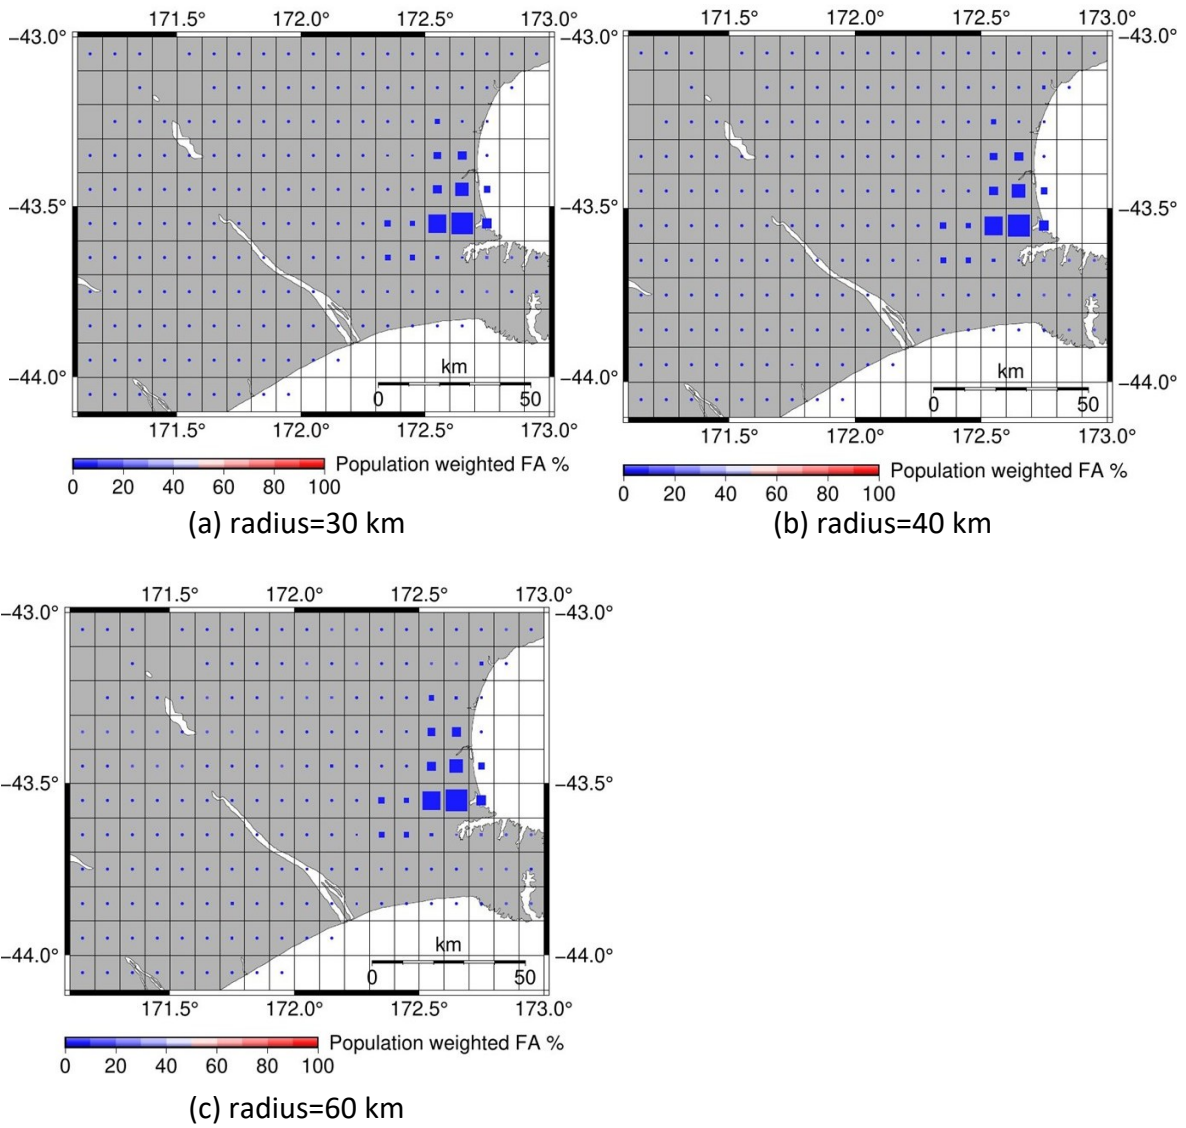


Figure 4-25 Population-weighted  $FA_{norm,g}$ , PLUM radius 30, 40, 60 km, other parameters are as in Figure 4-24.

Similarly, Figure 4-26 illustrates the Population-weighted normalized precautionary alerts for each grid ( $PA_{norm,g}$ ). The results show that precautionary alerts are more frequent in highly populated areas.

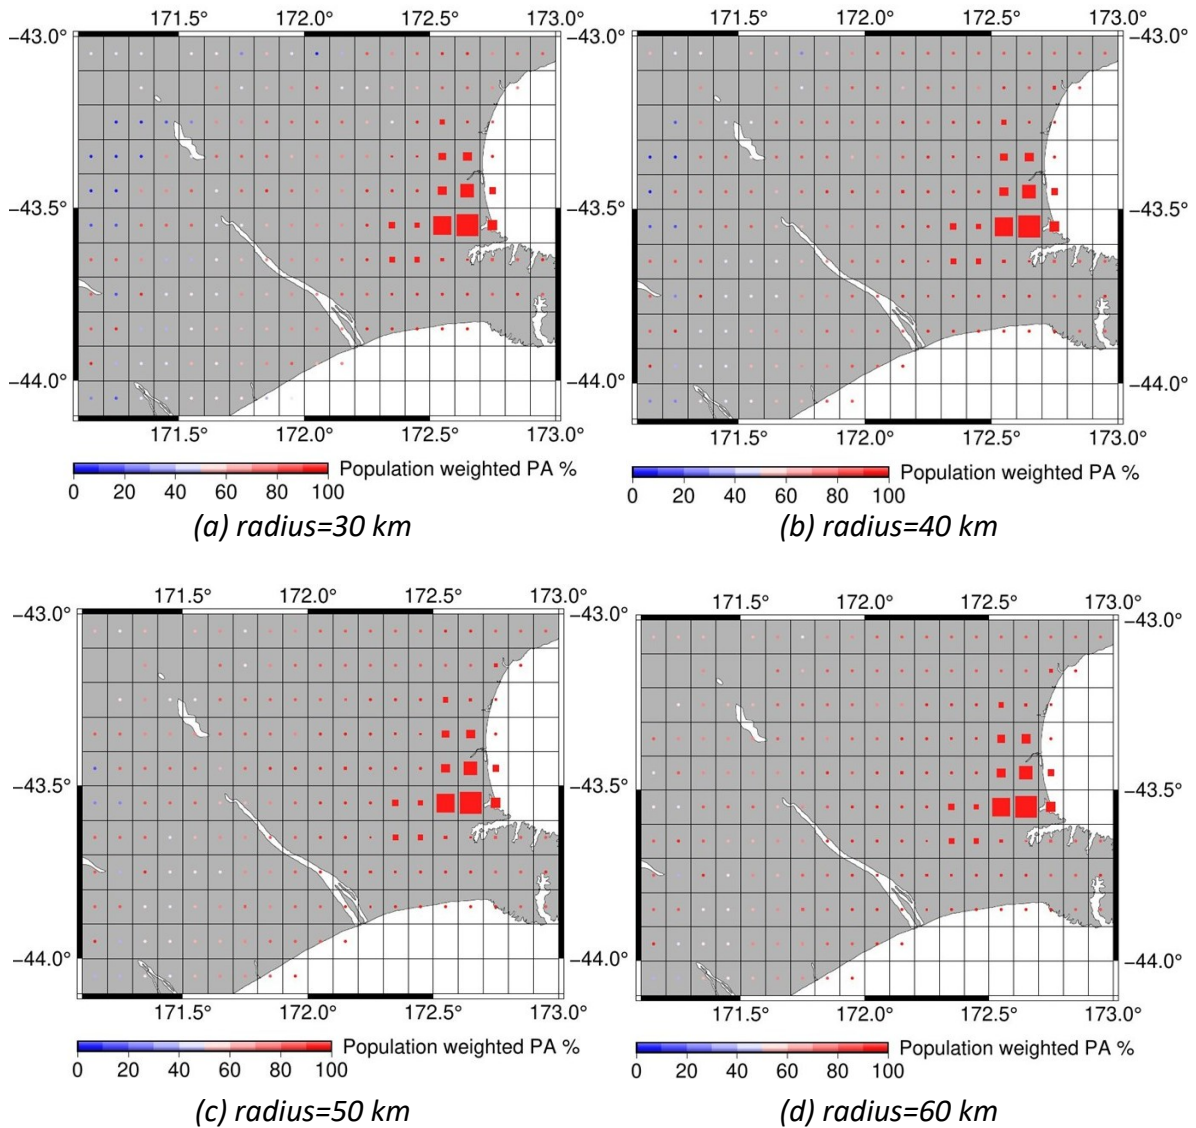


Figure 4-26 Population-weighted  $PA_{norm,g}$ : Precautionary Alerts vs. All Alerts for all 100 random 100-year periods (total 10,000 years), PLUM radius = (a) 30 km, (b) 40 km, (c) 50 km and (d) 60 km.  $Th=0.42 \text{ m/s}^2$ ,  $Tl=0.13 \text{ m/s}^2$ . The area of each square represents the number of person-alerts—the product of the population and the expected number of alerts received at that grid location (largest square=1853586, smallest square=6600). Squares that would have an area less than 1% of the largest square are shown as small fixed-size circles (less than 6,600 person-alerts in this case).

## Population-weighted overall region measures

Figure 4-22 to Figure 4-26 provide a useful visual representation of the geographically-distributed population-weighted Correct Timely Alerts, False Alerts and Precautionary Alerts. To give a measure of the performance over the whole region in order to facilitate comparisons across various parameters and settings, such as different PLUM warning radii and when adding or removing sensors, we use the population-weighted measures defined in equations 4.16 – 4.20 of Section 4.2.8.

The expected population-weighted  $CTA_{norm}$  score is measured by summing the CTA of each grid cell weighted by the population of the grid cell and then dividing by the sum of the product of the population of each cell and the total of the number of strong shaking events in the cell, as defined in equation 4.16. This gives the probability that an arbitrary person in the region who has experienced a strong shaking event will have received a correct timely alert for that event. The same approach applies to the  $FA_{norm}$  score and  $PA_{norm}$  score, as defined in equations 4.17 and 4.18. The population-weighted performance scores for various PLUM warning radii are presented in Table 4-8, along with the unweighted scores copied from Table 4-6 in Section 4.3.1. As in the unweighted case, the  $CTA_{norm}$  performance improves with increasing warning radius while the  $FA_{norm}$  performances get worse. However, the population-weighted  $PA_{norm}$  performance does not change much with warning radius.

*Table 4-8 Unweighted and Population-weighted  $CTA_{norm}$ ,  $FA_{norm}$ ,  $PA_{norm}$  for warning radii 30 km, 40 km, 50 km, 60 km.*

Warning radius	$CTA_{norm}$		$FA_{norm}$		$PA_{norm}$	
	Unweighted	Population weighted	Unweighted	Population weighted	Unweighted	Population weighted
30 km	54.7%	<b>89.6%</b>	1.1%	<b>0.2%</b>	80.3%	<b>94.9%</b>
40 km	68.1%	<b>92.6%</b>	1.9%	<b>0.2%</b>	84.9%	<b>95.1%</b>
50 km	74.4%	<b>94.1%</b>	3.3%	<b>0.4%</b>	86.0%	<b>95.1%</b>
60 km	79.1%	<b>94.8%</b>	4.9%	<b>0.8%</b>	86.1%	<b>94.8%</b>

Table 4-8 shows that the population-weighted scores for  $CTA_{norm}$  and  $FA_{norm}$  are better than the unweighted results. For example, the population-weighted  $CTA_{norm}$  performance score for the 50 km case (corresponding to the results shown in Figure 4-22) is 94.1% which is about 20% above the unweighted score of 74.4%. For the 30 km case, the

population-weighted  $CTA_{norm}$  is about 35% above the unweighted score. The population-weighted  $FA_{norm}$  is also better than the unweighted results, by a factor of more than 5. However, the population-weighted  $PA_{norm}$  is worse than the unweighted performance, increasing to about 95% for all warning radii.

The overall population-weighted  $CTA_{norm}$  and  $FA_{norm}$  results show that PLUM works well for the highly populated area of the Canterbury, though the high value of  $PA_{norm}$  is of concern. The performance can be attributed to the dense sensor network around Christchurch City, which also corresponds to areas of high population density.

### **Population-weighted warning time measures**

Section 4.3.1 discussed the warning times without considering the population. The expected mean warning times for correct alerts were illustrated in Figure 4-19; the same results are shown in Figure 4-27 and Figure 4-28, except that the sizes of the coloured squares are now determined by the product of population and the number of correct alerts. The size of a square can be interpreted as the expected number of times any person in the grid cell will receive a correct alert.

The expected mean warning time for the largest square in Figure 4-27 is 14 seconds for the 50 km warning radius. This is for the grid location within Christchurch City at latitude -43.55 and longitude 172.65, and is the most populated grid cell. For the second largest square, the expected mean warning time is 12.7 seconds.

The expected mean warning times for the same largest square in Figure 4-28 (Christchurch City) are 11.6 seconds, 13.2 seconds, and 14.5 seconds for the 30 km, 40 km, and 60 km respectively.

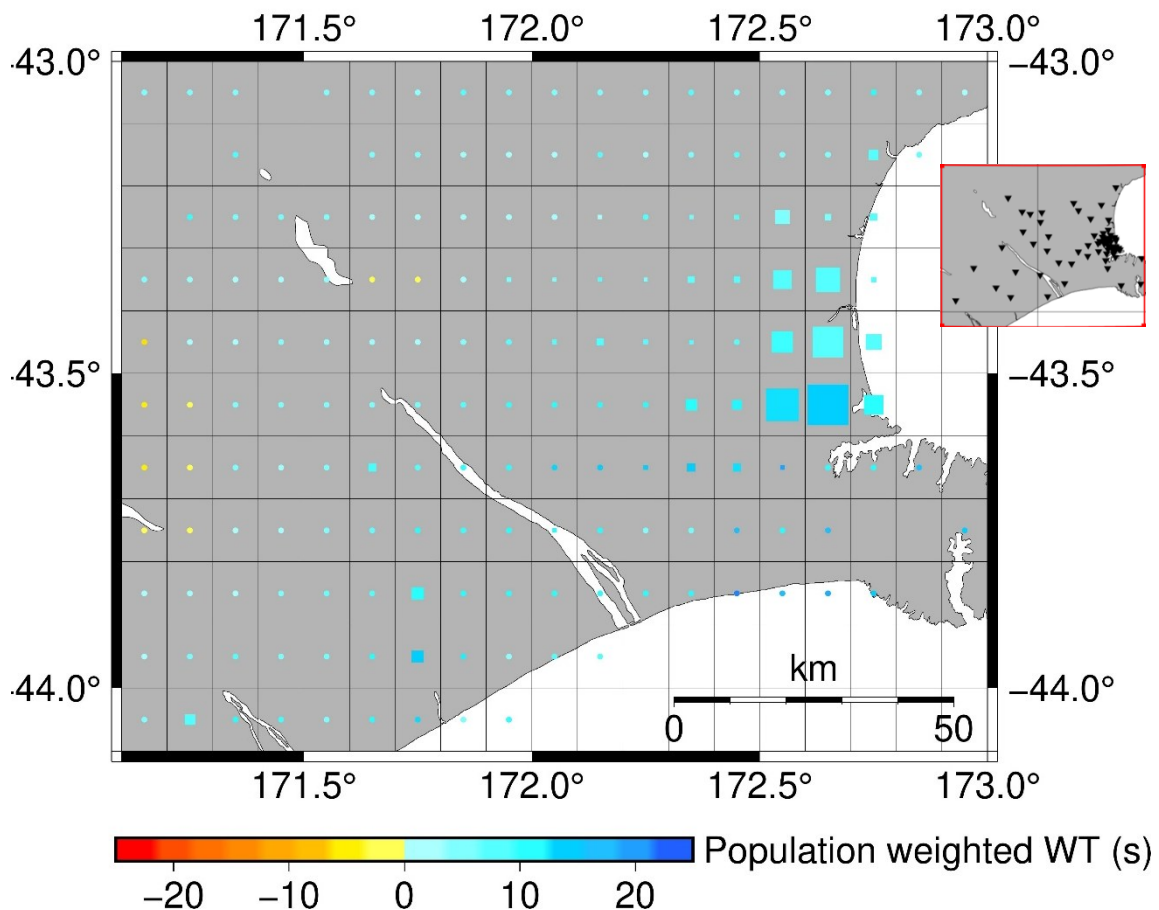


Figure 4-27 Population-weighted expected mean warning times for all correct alerts,  $WTCAG$ , for warning radius 50 km. Size of squares represents the expected number of correct person-alerts in the cell (largest square=2515581, smallest square=25688). Colour of squares represents the expected mean warning time: yellow for a negative mean warning time (late); blue for a positive mean warning time (timely). Squares that would have an area less than 1% of the largest square are shown as small fixed-size circles. ). Insert shows the positions of the sensors in the region (see Figure 4-3 for all sensors).

The overall population-weighted mean warning times for the region are shown in Table 4-9, along with the unweighted mean warning times copied from Table 4-7. The population-weighted warning times for all correct alerts are better than the unweighted warning times by almost 1.5 seconds. The population-weighted warning times for the correct **timely** alerts are better than the unweighted warning times by about 5 seconds, and also have a smaller variance. Combined with the population-weighted  $CTA_{norm}$  data in Table 4-8, this shows that with a warning radius of 50 km, a large fraction of the population of Canterbury would get around 11 seconds warning time for about 95% of strong shaking events.

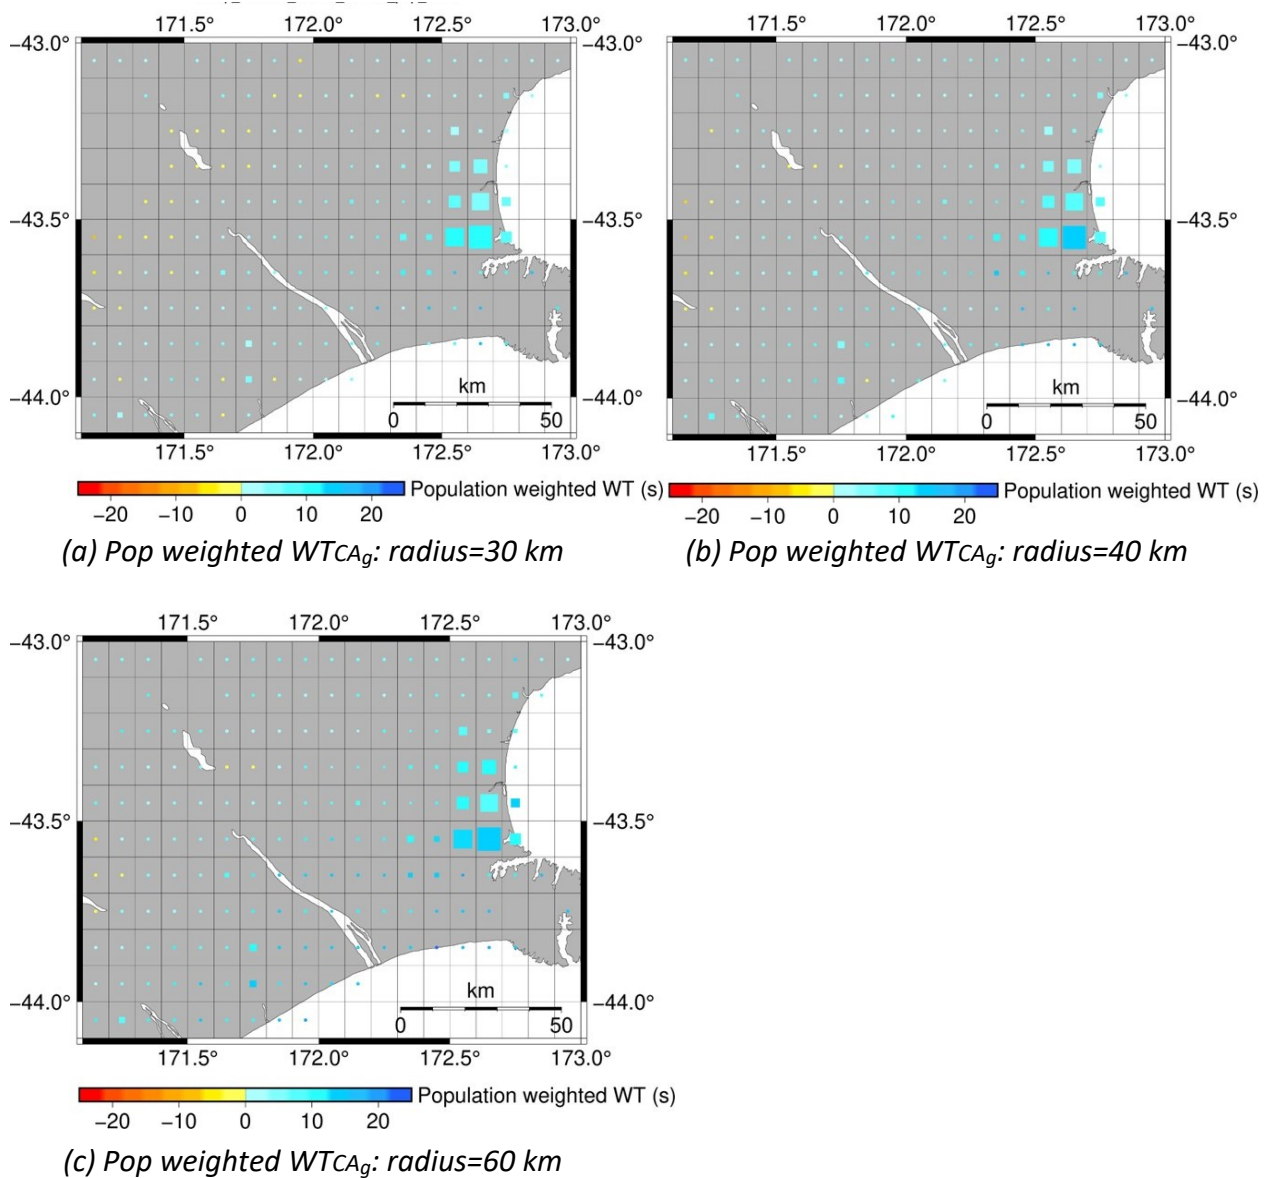


Figure 4-28 Population-weighted expected mean warning times for all correct alerts, for warning radii (a) 30 km, (b) 40 km and (c) 60 km. Size of squares represents the expected number of correct person-alert events in the cell (largest square=2515581, small square=25688). Colour of squares represents the expected mean warning time: yellow for a negative mean warning time (late); blue for a positive mean warning time (timely). Squares that would have an area less than 1% of the largest square are shown as small circles.

Table 4-9 Unweighted and Population-weighted mean warning times ( $\pm$  standard deviation) for all correct alerts, and for correct timely alerts, for warning radii 30 km, 40 km, 50 km, 60 km.

Warning radius	Mean WTCA (s)		Mean WTCTA (s)	
	Unweighted	Population weighted	Unweighted	Population weighted
30 km	2.7 $\pm$ 6.3	<b>7.5 <math>\pm</math> 3.4</b>	6.0 $\pm$ 4.3	<b>8.4 <math>\pm</math> 2.6</b>
40 km	4.5 $\pm$ 7.2	<b>9.0 <math>\pm</math> 3.2</b>	7.7 $\pm$ 5.1	<b>9.8 <math>\pm</math> 2.7</b>
50 km	6.0 $\pm$ 7.8	<b>10.3 <math>\pm</math> 3.1</b>	9.0 $\pm$ 5.8	<b>11.0 <math>\pm</math> 2.6</b>
60 km	7.4 $\pm$ 8.2	<b>11.4 <math>\pm</math> 2.6</b>	10.2 $\pm$ 6.5	<b>12.1 <math>\pm</math> 2.1</b>

### 4.3.3 Results for failure scenarios

The performance of the PLUM algorithm was also evaluated under two sensor failure scenarios, as detailed in Section 4.2.9. In the first scenario, a communication breakdown was simulated by randomly removing one-third of the 131 stations, representing a failure of one network service provider. In the second failure scenario, all stations within a 30 km to 60 km band around Christchurch were removed, representing a much less likely, but potentially more damaging scenario.

The results presented in this section follow the same format as Section 4.3.1 and 4.3.2. However, the graphs are provided only for the 50 km warning radius as the results for other warning radii show a very similar pattern of performance changes as the 50 km warning radius.

#### Removing 33% stations: grid cell measures

Figure 4-29 compares the expected normalized Correct Timely Alerts for each grid ( $CTA_{norm,g}$ ) for a 50 km warning radius under two scenarios (a) with all stations present and (b) with 33% of the stations removed. The right figure (b) shows the impact of removing 33% of the stations, where a small decrease in the percentage of correct timely alerts can be observed in some places, but there is no a major change in the performance.

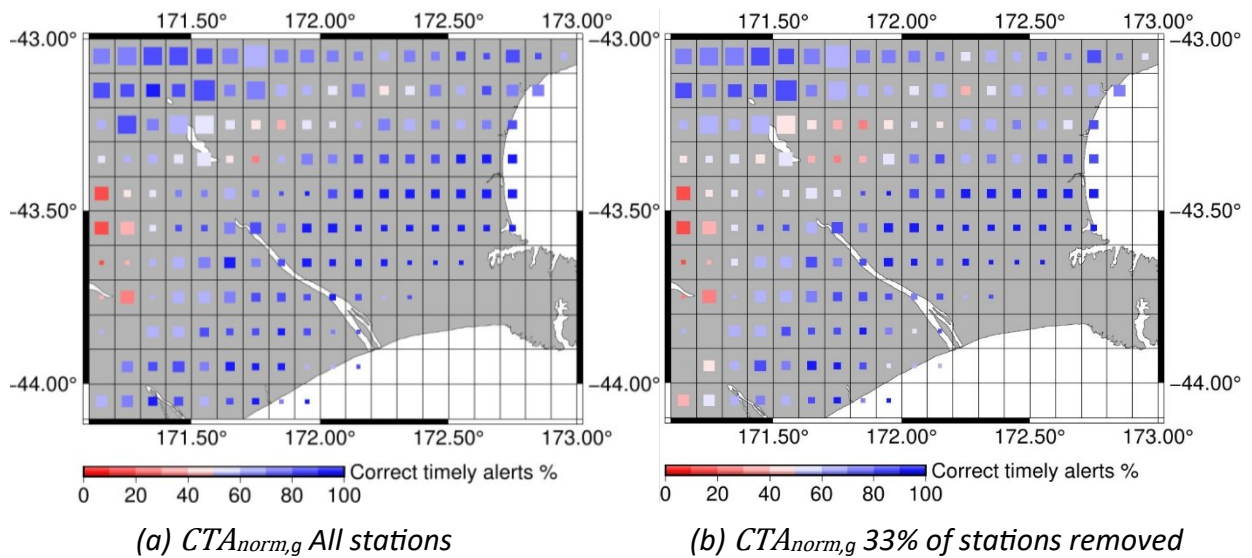


Figure 4-29  $CTA_{norm,g}$  for warning radius 50 km, with 33% of stations removed compared to all stations present. Size of squares represents number of strong shaking events in each grid cell (smallest square=0.1, largest square=2). Colour of squares represents the number of correct timely alerts as a fraction of the shaking events.

Figure 4

Figure 4-30 illustrates the expected normalized False Alerts for each grid cell ( $FA_{norm,g}$ ) for a 50 km warning radius under two scenarios: (a) with all stations present, and (b) with 33% of the stations removed. Comparing the two scenarios, removing 33% of the stations slightly decreases the percentage of false alerts, but does not significantly change across most grid cells, as indicated by the consistent blue shading.

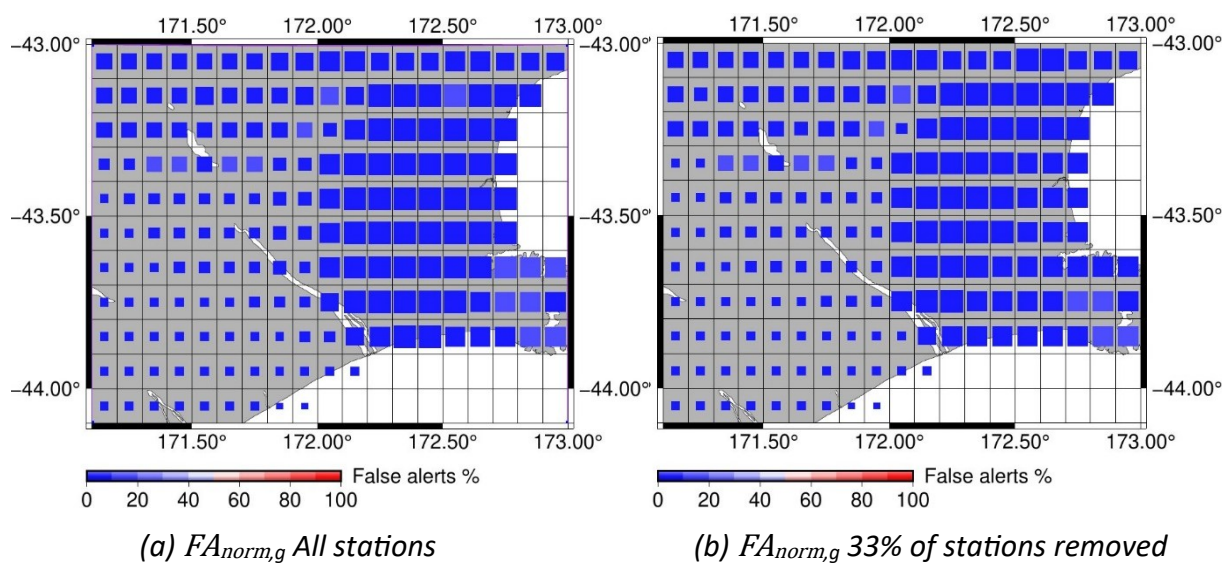


Figure 4-30  $FA_{norm,g}$  for warning radius 50 km, with 33% of the stations removed compared to all stations present. Size of squares represents the number of alerts received by each grid cell (smallest square= 0.6, largest square = 6.2) . Colour of squares represents the number of false alerts as a fraction of the total alerts.

Figure 4-31 displays the expected normalized Precautionary Alerts for each grid cell ( $PA_{norm,g}$ ) for a 50 km warning radius under two scenarios: (a) with all stations present, and (b) with 33% of the stations removed. The comparison between the two scenarios indicates that removing 33% of the stations results in a slight decrease in percentage of Precautionary Alerts in some locations but does not significantly change across most grid cells.

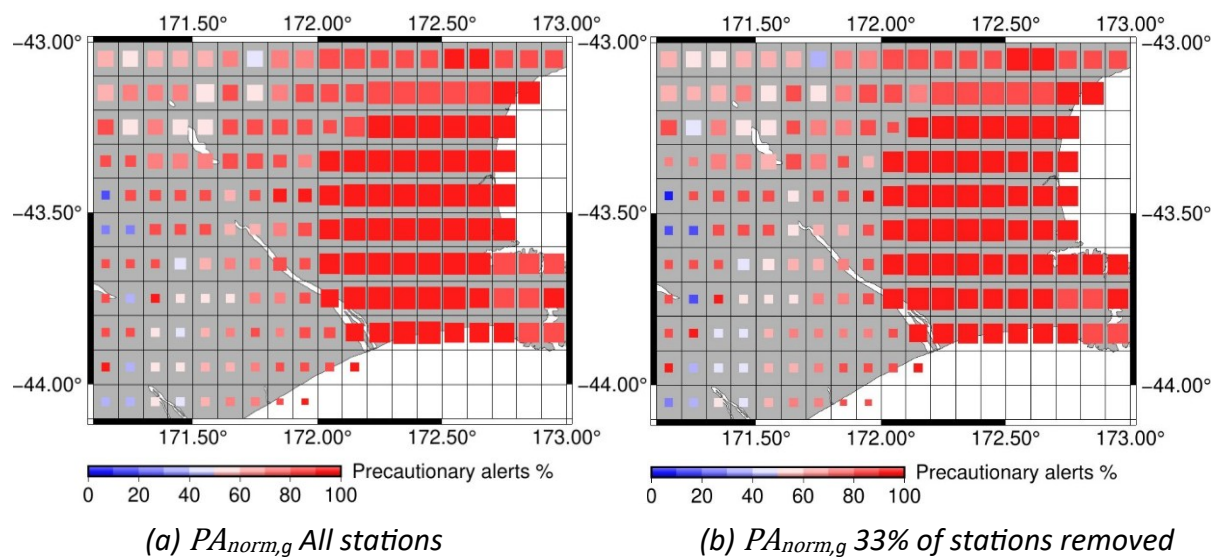


Figure 4-31  $PA_{norm,g}$  for warning radius 50 km, with 33% of the stations removed compared to all stations present. Size of squares represents the number of alerts received by each grid cell (smallest square= 0.6, largest square = 6.2). Colour of squares represents the number of precautionary alerts as a fraction of the total alerts.

### Removing 33% stations: region measures

Table 4-10

Table 4-10 presents the expected values of  $CTA_{norm}$ ,  $FA_{norm}$  and  $PA_{norm}$  for the entire region for different warning radii comparing scenarios with all stations operational and with 33% of the stations removed. The table shows that removing 33% of the stations reduces the  $CTA_{norm}$  probability across all warning radii, with the most significant drop at the 30 km radius (from 54.7% to 46.9%). But the reduction for the 50 km radius is 5.8% only. The results indicate a slight decrease in  $FA_{norm}$  with station removal, suggesting a small reduction in the false alerts.  $PA_{norm}$  values also show a slight decrease.

Table 4-10 Expected  $CTA_{norm}$ ,  $FA_{norm}$  and  $PA_{norm}$  with 33% of the stations removed compared with all stations working, across whole region and for warning radii 30 km, 40 km, 50 km, 60 km.

Warning radius	Expected $CTA_{norm}$		Expected $FA_{norm}$		Expected $PA_{norm}$	
	All stations	33% removed	All stations	33% removed	All stations	33% removed
30 km	54.7%	46.9%	1.1%	0.9%	80.3%	78.6%
40 km	68.1%	60.8%	1.9%	1.8%	84.9%	84.0%
50 km	74.4%	68.6%	3.3%	3.0%	86.0%	85.5%
60 km	79.1%	74.0%	4.9%	4.7%	86.1%	85.8%

Table 4-11 Table 4-11 presents the expected number of each alert outcome for the region for different warning radii comparing scenarios with all stations operational and with 33% of the stations removed. The results show that removing 33% of the stations generally leads to a slight decrease in the number of CTA across all warning radii. There is an increase in LA and MA, indicating that some CTA alerts become either late or missed when fewer stations are available. For PA and FA, the results suggest a decrease in the number of precautionary and false alerts when stations are removed. But the changes are minor, showing that the PLUM algorithm can handle station loss well.

Table 4-11 Expected number of each PLUM outcome with 33% of the stations removed compared with all stations working, across whole region and for warning radii 30 km - 60 km.

Warning radius	CTA		LA		MA		MLA		PA		FA	
	All	33%	All	33%	All	33%	All	33%	All	33%	All	33%
30 km	37	32	17	18	13	17	31	36	291	259	4	3
40 km	46	41	16	18	6	9	21	26	434	398	10	8
50 km	50	46	15	17	2	4	17	21	540	504	20	18
60 km	53	50	13	15	1	2	14	18	642	609	36	33

Table 4-12 Table 4-12 shows the population-weighted values for  $CTA_{norm}$ ,  $FA_{norm}$  and  $PA_{norm}$  with all stations operational and with 33% of the stations removed. The population-weighted  $CTA_{norm}$ , slightly decreases when 33% of the stations are removed. The population-weighted  $FA_{norm}$  values remain very low and unchanged at most warning radii except 60 km, even after station removal.  $PA_{norm}$  shows a negligible decrease with station removal as it got high values in both cases.

Table 4-12 Population-weighted  $CTA_{norm}$ ,  $FA_{norm}$ ,  $PA_{norm}$  with 33% of the stations removed compared with all stations working, across whole region and for warning radii 30 km, 40 km, 50 km, 60 km.

Warning radius	Population-weighted $CTA_{norm}$		Population-weighted $FA_{norm}$		Population-weighted $PA_{norm}$	
	All stations	33% removed	All stations	33% removed	All stations	33% removed
30 km	89.6%	85.1%	0.2%	0.2%	94.9%	94.8%
40 km	92.6%	89.1%	0.2%	0.2%	95.1%	94.9%
50 km	94.1%	90.8%	0.4%	0.4%	95.1%	94.9%
60 km	94.8%	92.0%	0.8%	0.7%	94.8%	94.7%

### Removing 33% stations: warning time

Figure 4-32 Figure 4-32 illustrates the expected mean warning times for all correct alerts within a 50 km radius comparing two scenarios: (a) when all stations are included, and (b) when 33% of the stations are removed. As shown in the figure, even after removing a significant number of stations, the overall pattern of warning times stays mostly the same with only a very little reduction in the scenario in which 33% of stations are removed.

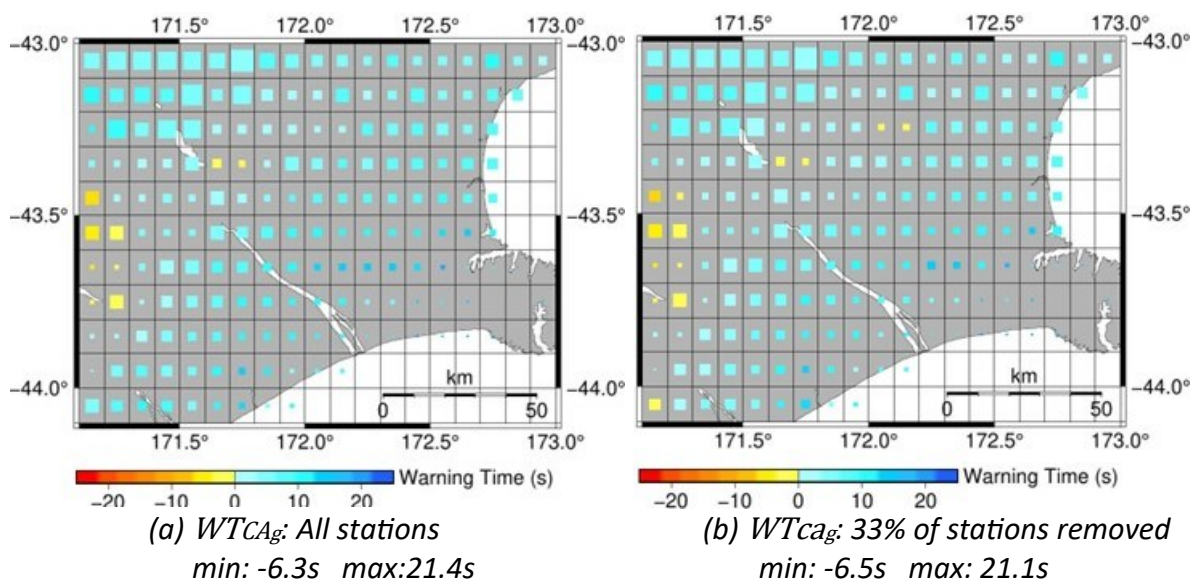


Figure 4-32 Expected  $WT_{cag}$  with 33% of stations removed for warning radius 50 km, compared with all stations. Size of squares represents the expected number of correct alerts in the cell (small square=0.1, largest square=1.8). Colour of squares represents the expected mean warning time: yellow/orange for a negative mean warning time (late); blue for a positive mean warning time (timely).

Figure 4-33 Figure 4-33 presents the warning times distribution of all correct alerts across the region, displayed as box plots. The plots compare scenarios with all stations (blue) and with 33% of the stations removed (orange). Each box plot displays the median (red line), along with the 25th and 75th percentiles, and includes outliers to show the range of variation in warning times. It is clear that distribution is slightly impacted by the removal of station but there is no significant change.




Figure 4-33 Distribution of  $WT_{CA}$  with 33% of stations removed: box plots of warning times of all correct alerts across region, with all stations (blue), and with 33% of the stations removed (orange). Warning radii 30 km, 40 km, 50 km, 60 km. The box plots show the median (red line), 25<sup>th</sup> and 75<sup>th</sup> percentiles, and the outliers.

Figure 4-34 Figure 4-34 illustrates the distribution of warning times for all correct **timely** alerts across the region to compare scenarios with all stations included (blue) versus 33% of the stations removed (orange). The results indicate that while the removal of 33% of stations does slightly impact the distribution of warning times, the overall effect is not significant.

Table 4-13 Table 4-13 shows the mean warning times along with their standard deviations, both before and after the removal of stations. The removal of 33% of the stations results in a

reduction of the mean warning time for correct timely alerts by up to 0.5 seconds, and for all correct alerts by up to 1 second. There is very little change to the standard deviation. It is showing that the warning times are still mostly consistent even after reducing the number of stations.

Based on all results of the 33% random station removal scenario, we conclude that the PLUM EEW performance remains relatively stable even with a reduction in the number of stations.




Figure 4-34 Distributions of *WTCTA* with 33% of stations removed: box plots of warning times of all correct **timely** alerts across region with all stations (blue), and with 33% of stations removed (orange). Warning radii 30 km, 40 km, 50 km, 60 km. The box plots show the median (red line), 25<sup>th</sup> and 75<sup>th</sup> percentiles, and the outliers.

Table 4-13 Mean  $\pm$  standard deviation of distributions of *WTCTA* and *WTCA* (warning time for timely alerts and for all correct alerts) with 33% of stations removed and with all stations working, for different warning radii.

Warning Radius	WTCTA (s)		WTCA (s)	
	All stations	33% removed	All stations	33% removed
30 km	6.0 $\pm$ 4.3	5.9 $\pm$ 4.4	2.7 $\pm$ 6.3	2.0 $\pm$ 6.6
40 km	7.7 $\pm$ 5.1	7.4 $\pm$ 5.0	4.5 $\pm$ 7.2	3.8 $\pm$ 7.3

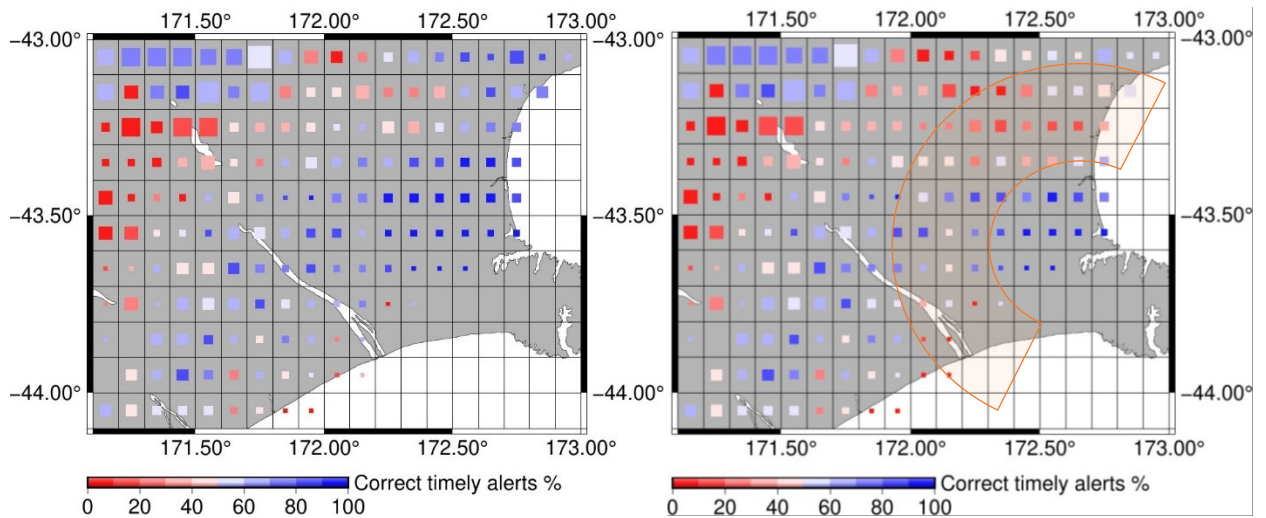
50 km	9.0 ± 5.8	8.6 ± 5.6	6.0 ± 7.8	5.0 ± 7.8
60 km	10.2 ± 6.5	9.7 ± 6.3	7.4 ± 8.2	6.4 ± 8.3

### Removing stations from band: grid cell measures

In the second failure scenario, all stations within a 30 km to 60 km band around Christchurch (centred at the grid cell with the highest population where Christchurch City is located) were removed (see Figure 4-12).

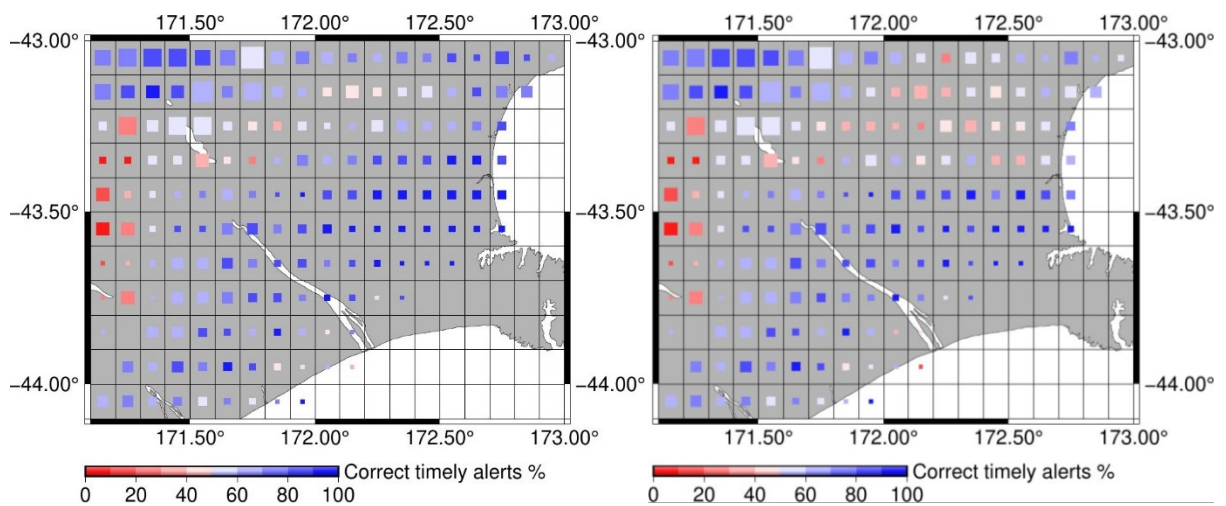
Figure 4-35 presents the expected  $CTA_{norm,g}$  results for each grid for each warning radius, comparing (on the left) the original scenario where all stations are present, and (on the right) the scenario where stations in this band were removed. Unlike the previous scenario, the results for all warning radii are shown because there are some visible differences between the different warning radii. Particularly at the lower warning radii, removing stations in the band around Christchurch visibly reduces the  $CTA_{norm,g}$  performance immediately to the north and west of Christchurch, though it does not appear to affect performance in Christchurch itself very much. At 60 km warning radius, removing the stations appears to have little effect, except in the northern part of the band itself.

Figure 4-36 and Figure 4-37 illustrate the expected  $FA_{norm,g}$  and  $PA_{norm,g}$  results for the 50 km warning radius, comparing (a) the results with all stations present and (b) the results after the removal of stations in the specified band around Christchurch. For both  $FA_{norm,g}$  and  $PA_{norm,g}$ , the figures show that the removal of stations in the band did not have a visible impact on the false alert and precautionary alert ratios. The results for other warning radii are not shown because that are very similar to the results for the 50 km warning radius.



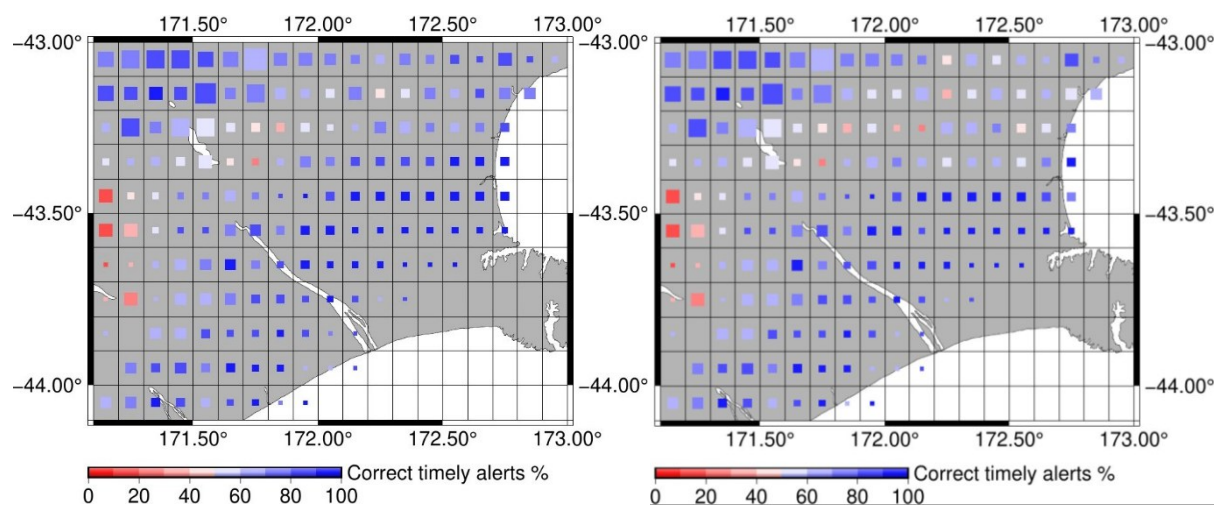
(a)  $CTA_{norm,g}$  30 km, all stations

(b)  $CTA_{norm,g}$  30 km, stations in band removed



(c)  $CTA_{norm,g}$  40 km, all stations

(d)  $CTA_{norm,g}$  40 km, stations in band removed



(e)  $CTA_{norm,g}$  50 km, all stations

(f)  $CTA_{norm,g}$  50 km, stations in band removed

(figure continued on next page)

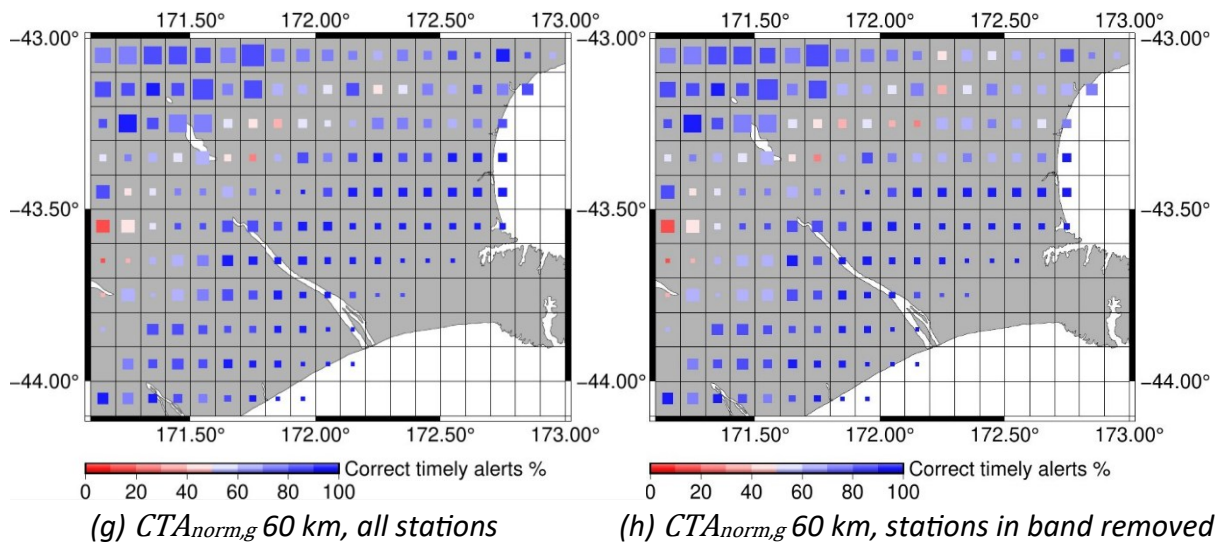


Figure 4-35  $CTA_{norm,g}$  with stations in band removed: (a),(c),(e),(g):  $CTA_{norm,g}$  with all stations present, and (b),(d),(f),(h):  $CTA_{norm,g}$  with stations in band 30-60 km from Christchurch removed. For warning radii 30 km, 40 km, 50 km and 60 km, The orange highlighted region in (b) shows the band where stations are removed (also see Figure 4-12). Size of squares represents number of strong shaking events in each grid cell (smallest square=0.1, largest square=2). Colour of squares represents the number of correct timely alerts as a fraction of the shaking events.

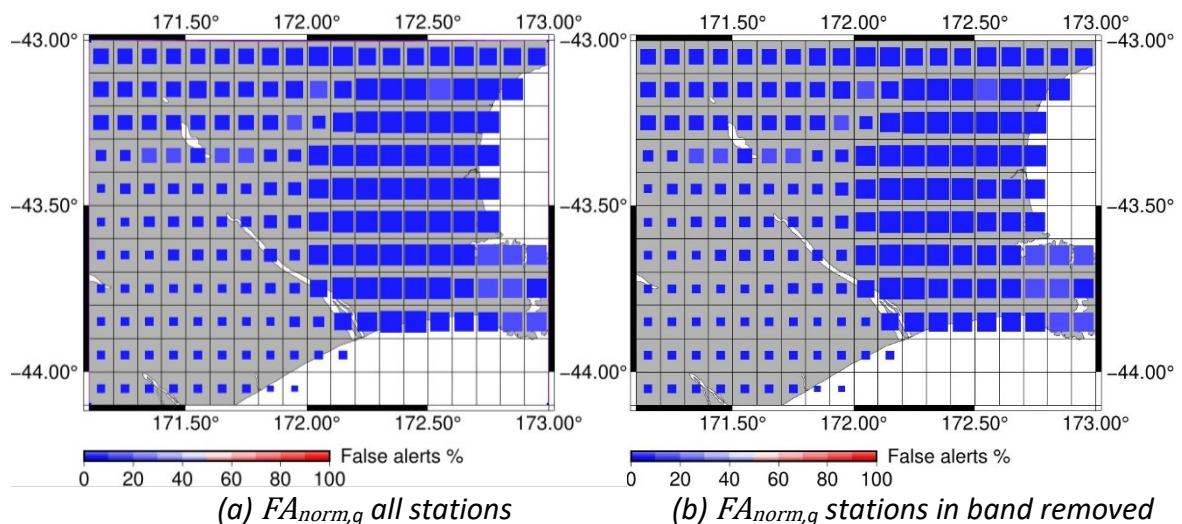


Figure 4-36  $FA_{norm,g}$  with stations in band removed for warning radius 50 km:  $FA_{norm,g}$  with (a) all stations present, and (b) stations in band 30-60 km from Christchurch removed. Size of squares represents the number of all alerts received by each grid cell (smallest square= 0.6, largest square = 6.2). Colour of squares represents the number of false alerts as a fraction of the total alerts.

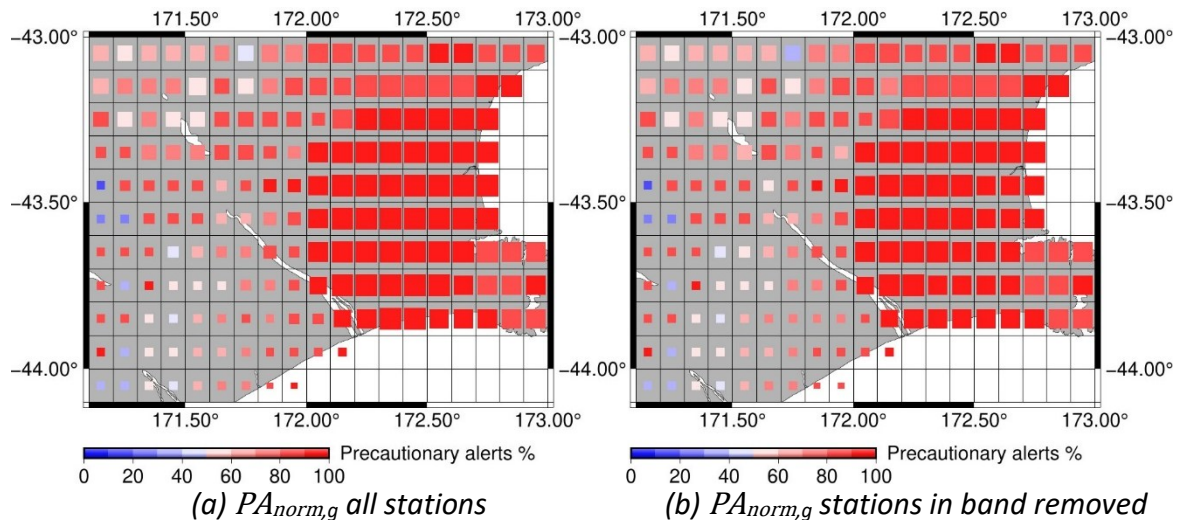


Figure 4-37  $PA_{norm,g}$  with stations in band removed for warning radius 50 km:  $PA_{norm,g}$  with (a) all stations present, and (b) all stations in band 30-60 km from Christchurch removed. Size of squares represents the number of all alerts received by each grid cell (smallest square= 0.6, largest square = 6.2). Colour of squares represents the number of precautionary alerts as a fraction of the total alerts.

#### Removing stations from band: region measures

Table 4-14 Table 4-14 presents the expected values of  $CTA_{norm}$ ,  $FA_{norm}$  and  $PA_{norm}$  across the entire region for different warning radii, comparing scenarios with all stations operational and with stations in a 30-60 km band around Christchurch removed. The results indicate that while the removal of stations in this band does have some effect on the performance, the impact is relatively tolerable. For instance, the expected  $CTA_{norm}$  decreases slightly across all warning radii with the largest drop of 7% at the 30 km warning radius (from 54.7% to 47.7%) and the smallest drop of only 2.5% for the 60 km warning radius. The expected  $FA_{norm}$  shows almost no change with the removal of the band stations. Similarly, the expected  $PA_{norm}$  indicates only a minor decrease when the stations are removed.

Table 4-14 Expected  $CTA_{norm}$ ,  $FA_{norm}$  and  $PA_{norm}$ , with station in band removed compared with all stations working, across whole region and for different warning radii. Band covers stations 30-60 km from Christchurch.

Warning radius	Expected $CTA_{norm}$		Expected $FA_{norm}$		Expected $PA_{norm}$	
	All stations	Band removed	All stations	Band removed	All stations	Band removed
30 km	54.7%	47.7%	1.1%	1.1%	80.3%	79.5%
40 km	68.1%	63.2%	1.9%	1.9%	84.9%	84.5%
50 km	74.4%	71.3%	3.3%	3.2%	86.0%	85.8%
60 km	79.1%	76.6%	4.9%	4.8%	86.1%	86.0%

Table 4-15 shows the expected number of each alert outcome across the entire region for different warning radii comparing scenarios with all stations operational and with stations in a 30-60 km band around Christchurch removed. There is a slight reduction in the number of CTA when the band of stations is removed, particularly at 30 km warning radius (a decrease from 37 to 32). But the decrease becomes less noticeable as the warning radius increases. The number of LA slightly increases when stations in the band are removed, presumably because some alerts are being issued by stations closer to the receiving grid point. For example, at a 30 km radius, LA increase from 17 to 20. MA also shows a small increase when stations are removed. The number of PA decreases slightly when the band of stations is removed, and FA almost remains unchanged.

Table 4-15 Expected number of each PLUM outcome with station in band removed compared with all stations working, across whole region and for different warning radii. Band covers stations 30-60 km from Christchurch.

Warning radius	CTA		LA		MA		MLA		PA		FA	
	All	Band	All	Band	All	Band	All	Band	All	Band	All	Band
30 km	37	32	17	20	13	15	31	35	291	277	4	4
40 km	46	43	16	18	6	6	21	25	434	421	10	10
50 km	50	48	15	17	2	3	17	19	540	529	20	20
60 km	53	52	13	14	1	1	14	16	642	633	36	35

Table 4-16 presents the population-weighted values for  $CTA_{norm}$ ,  $FA_{norm}$  and  $PA_{norm}$  across different warning radii comparing scenarios where all stations are operational versus when stations in a 30-60 km band around Christchurch are removed. There is a noticeable decrease in the population-weighted  $CTA_{norm}$  when the band of stations is

removed, particularly at the 30 km radius (dropping from 89.6% to 78.2%). However, this reduction decreases as the warning radius increases (e.g., only 94.8% to 90.2% at 60 km). The population-weighted  $FA_{norm}$  remains nearly unchanged across all radii with only a slight decrease observed at the 60 km radius (from 0.8% to 0.7%). The population-weighted  $PA_{norm}$  remains constant across all radii with no change observed when the band of stations is removed.

*Table 4-16 Population-weighted  $CTA_{norm}$ ,  $FA_{norm}$ ,  $PA_{norm}$  with station in band removed compared with all stations working, across whole region and for different warning radii. Band covers stations 30-60 km from Christchurch.*

Warning radius	Population-weighted $CTA_{norm}$		Population-weighted $FA_{norm}$		Population-weighted $PA_{norm}$	
	All stations	Band removed	All stations	Band removed	All stations	Band removed
30 km	89.6%	78.2%	0.2%	0.2%	94.9%	94.9%
40 km	92.6%	82.5%	0.2%	0.2%	95.1%	95.1%
50 km	94.1%	87.4%	0.4%	0.4%	95.1%	95.1%
60 km	94.8%	90.2%	0.8%	0.7%	94.8%	94.8%

#### **Removing stations from band: warning time**

Figure 4-38 displays the expected mean warning times for all correct alerts within 50 km warning radius, comparing two scenarios: (a) with all stations operational, and (b) with stations in the 30-60 km band around Christchurch removed. The comparison shows that removing stations in the 30-60 km band results in only minor changes to the overall warning times, with the maximum and minimum values remaining consistent between the two scenarios. The results are very similar at other warning radii and are therefore not shown.

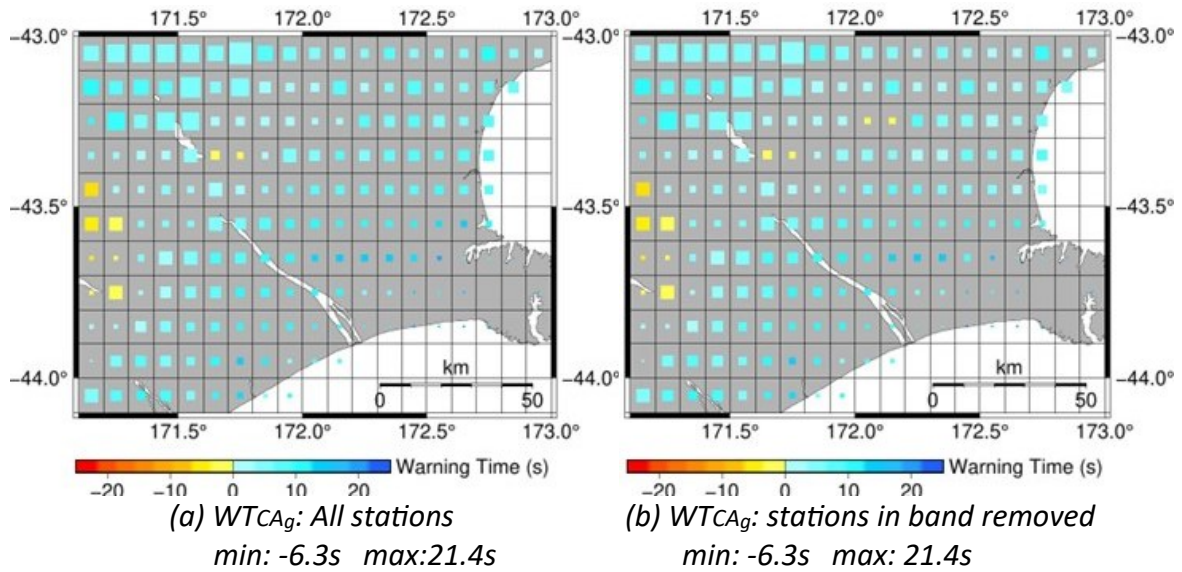


Figure 4-38 Expected  $WTCA_g$  with stations in band removed: mean warning times for all correct alerts for warning radius 50 km, (a) with all stations, and (b) with stations in band 30 km-60 km from Christchurch removed. Size of squares represents the expected number of correct alerts in the cell (small square=0.1, largest square=1.8). Colour of squares represents the expected mean warning time: yellow/orange for a negative mean warning time (late); blue for a positive mean warning time (timely).

Figure 4-39 shows the distributions of warning times for correct alerts through box plots. The figure contrasts two scenarios: one with all stations operational (blue) and another with the band of stations between 30 km and 60 km from Christchurch removed (orange). The comparison shows that removing the band of stations causes some changes in the distribution of warning times, but the overall pattern stays consistent.




Figure 4-39 Distributions of *WTCA* with band of stations removed: box plots of warning times of all correct alerts across region, with all stations (blue), and with band of stations between 30 km-60 km from Christchurch removed (orange). Warning radii 30 km, 40 km, 50 km, 60 km. The box plots show the median (red line), 25<sup>th</sup> and 75<sup>th</sup> percentiles, and the outliers.

Figure 4-40 shows the distributions of warning times for correct timely alerts to compare warning times across the region under both scenarios. As with all correct alerts above, the comparison indicates that while removing the band of stations causes some small changes in the distribution of warning times, the overall pattern remains consistent across the different warning radii.




Figure 4-40 Distributions of  $WT_{CTA}$  with band of stations removed : box plots of warning times of all correct alerts across region with all stations (blue), and with band of stations between 30 km-60 km from Christchurch removed (orange). Warning radii 30 km, 40 km, 50 km, 60 km. The box plots show the median (red line), 25<sup>th</sup> and 75<sup>th</sup> percentiles, and the outliers.

Table 4-17 Table 4-17 presents the mean and standard deviation of the distributions of  $WT_{CTA}$  (warning time for timely alerts only) and  $WT_{CA}$  (warning time for all correct alerts) across different warning radii. The table compares the scenarios with all stations operational and with the band of stations between 30 km and 60 km from Christchurch removed. The mean warning times for timely alerts slightly decrease when the band of stations is removed, but the change is minimal. For instance, at the 30 km radius, the mean  $WT_{CTA}$  drops slightly from 6.0 seconds with all stations to 5.9 seconds with the band removed. The standard deviation stays consistent, showing similar variability in warning times for both scenarios. The mean  $WT_{CA}$  also decreases when the band of stations is removed, with a larger reduction compared to  $WT_{CTA}$ . At the 30 km radius, the mean  $WT_{CA}$  drops from 2.7 seconds to 1.9 seconds. As the warning radius increases, the difference between the scenarios gets smaller indicating that removing the band of stations has less impact at larger radii.

Table 4-17 Mean  $\pm$  standard deviation of distributions of *WTCTA* and *WTCA* (warning time for timely alerts and for all correct alerts) with band of stations between 30 km-60 km from Christchurch removed and with all stations working, for different warning radii.

Warning Radius	WTCTA (s)		WTCA (s)	
	All stations	Band removed	All stations	Band removed
30 km	6.0 $\pm$ 4.3	5.9 $\pm$ 4.4	2.7 $\pm$ 6.3	1.9 $\pm$ 6.5
40 km	7.7 $\pm$ 5.1	7.4 $\pm$ 5.1	4.5 $\pm$ 7.2	3.8 $\pm$ 7.3
50 km	9.0 $\pm$ 5.8	8.8 $\pm$ 5.8	6.0 $\pm$ 7.8	5.4 $\pm$ 7.8
60 km	10.2 $\pm$ 6.5	10.1 $\pm$ 6.5	7.4 $\pm$ 8.2	7.0 $\pm$ 8.4

Overall, the results suggest that removing the band of stations between 30 km and 60 km from Christchurch results in a slight decrease in the performance. However, the changes are relatively small, and the PLUM EEW continues to provide effective EEW even with the loss of stations.

#### 4.4 Discussion

This section discusses the results from Section 4.3, focusing on issues that are specific to the Canterbury region. Issues that apply to both the Canterbury region and the Wellington region (Chapter 5) are discussed in Chapter 6.

The results presented in Section 4.3.1 suggest that Christchurch City and its surroundings (the central area of the eastern part of the map) almost always get correct timely alerts, whereas the west region of the area gets more missed and late alerts. According to Figure 4-14, the expected fraction of correct timely alerts (out of all strong shaking events) exceeds 50% for the majority of locations. However, a few locations, primarily situated in the western half of the region, have a higher fraction of missed or late alerts compared to correct timely alerts. Almost all of those poorly performing locations are in the western half of the region and the low station density around the western half is the major cause of this (see Figure 4-3). This outcome is expected as PLUM was originally designed for EEW in Japan which has a higher station density than New Zealand.

The results suggested the 30 km warning radius setting is not suitable for the NZ context given the density of GeoNet strong motion sensor stations. This suggests that expanding the radius to 40 km or 50 km would lead to significant improvements in the performance compared to

the 30 km setting. The gains become marginal when extending the radius from 50 km to 60 km. However, a higher warning radius also gets more false and precautionary alerts. Again, this outcome is expected, as the PLUM algorithm assumes undamped seismic motion, which is a reasonable assumption for a 30 km radius but becomes less accurate at larger distances.

According to the overall measures for the entire region as shown in Table 4-6, using a 50 km warning radius results in fairly good performance: when a strong shaking event occurs within the region, there is a 74% probability of a location receiving a correct timely alert rather than a missed or late alert and when an alert is received at a location, there is less than a 5% chance that the alert is a false alarm. On the other hand, the fraction of alerts that are precautionary alerts is always greater than 80% for all 30 to 60 km warning radius settings.

However, to better reflect the risk of possible injury and damage, it is necessary to take the population into account, and Section 4.3.2 presents the results taking the population of each grid cell in the region into consideration. The population-based performance analysis presented in Section 4.3.2 indicates that the PLUM EEW algorithm performs well in the high population areas of the Canterbury region. The main reason for that is the high station density around the high population areas including Christchurch City. Overall, for the 50 km warning radius scenario, the normalised correct timely alert value improves from 74% to 94% when the population is taken into account as shown in Table 4-8, and the false alert fraction reduces to 0.4%. However, the precautionary alert ratio is still high, which requires further discussion.

Researchers have shown that repeated false or precautionary alerts can lead to a phenomenon known as "alert fatigue," or "warning fatigue" where people start to ignore warnings because they experience them too often without significant consequences (Kolb & Conitzer, 2020; Mackie, 2014). While these studies are from other fields and not directly related to EEW, the concept of warning fatigue remains relevant. A high fraction of precautionary alerts is therefore of concern. However, according to the synthetic catalogue data used in this analysis, the Canterbury region can expect approximately 10 large earthquakes over a 100-year period. This translates to about one EEW alert every 10 years. Even during more active periods, the region is unlikely to receive more than 10 alerts per decade, averaging around one per year. Because of this low frequency, alert fatigue seems unlikely, and people probably would not ignore the alerts, even if many of them turn out to be precautionary alerts.

Nonetheless, if there are many aftershocks after a large event, the higher frequency of precautionary alerts may cause problems; Becker et al. (2020a) mentioned that people might start ignoring warnings during an aftershock sequence if they receive too many alerts. Therefore, how to deal with a precautionary alert during an aftershock sequence to avoid alert fatigue needs further research. Changing the alert threshold to a higher value during an aftershock sequence may be a solution to avoid alert fatigue in this particular situation.

In practice, PLUM would generate more precautionary and false alerts than in the analysis presented here, because real earthquakes include many smaller events that still exceed trigger thresholds. In operational mode, the thresholds used for detection and confirmation stations should be adjusted to filter out some of those false activations from small events. Thus, any operational PLUM EEW system must account for a higher rate of small-event activations when forecasting false-alert fractions.

Another possible cause of alert fatigue was mentioned by Vinnell et al. (2023) who state that setting the alert threshold too low can cause people to ignore frequent warnings and lead to warning fatigue. However, the threshold used in this study is MMI V, which is higher than the thresholds of some other EEW systems such as the US Shake Alert or the threshold used in the West Coast of US testing of PLUM (Cochran et al., 2022).

According to Figure 4-18 of Section 4.3.1, expected mean warning times from the alerts are always positive for Christchurch City and surrounding grid locations (at best, 22 seconds with a 60 km warning radius) although some locations in the rural western part of the region have mostly late alerts, with a negative expected mean warning times (at worst -8 seconds with a 60 km warning radius). According to Figure 4-19 and Table 4-7, the mean warning time of correct timely alerts over the whole region for the 50 km warning radius is 9 seconds. When weighted by population, Figure 4-26 and Table 4-9 show that the highly populated areas always get positive warning times (i.e., timely alerts) and that the mean warning time rises by about 2 seconds, (11 seconds in the 50 km case). Therefore, the PLUM algorithm could provide many effective alerts with sufficient warning time to take immediate actions. However, the warning times are not uniform, and the distributions show that for some of the correct timely alerts, the warning time may be only a few seconds, which may not allow people to take effective preventative actions.

Results presented in Sections 4.3.1 and 4.3.2 assumed an ideal scenario where the EEW system operates without any interruptions. However, infrastructure could fail during an earthquake which may impact the system's performance. Section 4.3.3 presented results for PLUM in two sensor network failure scenarios. The first scenario examines the impact of losing data from 33% of the stations, which could simulate the failure of one network provider or random sensor outages. The analysis shows that the overall system performance is minimally affected, even with a 33% station loss. The expected fraction of correct timely alerts for the region decreases by only 5% for the 50 km radius case (from 74% to 69%, as shown in

Table 4-10) and warning times decreased by less than 1 second (see Table 4-13) For the population-weighted results, the correct timely alert ratio reduction is even smaller—just 3% with the 50 km radius (from 94% to 91%, according to Table 4-12). This suggests that the PLUM algorithm could work effectively even with the random station outages in the Canterbury region without significantly affecting the performance.

The second failure scenario involves removing sensors within a band 30 km to 60 km from Christchurch City. The results show that this didn't badly affect performance, especially for the 50 and 60 km warning radii, with the reduction in the overall correct timely alert ratios for the region (as shown in Table 4-14) being less than in the 33% station removal case. However, the population-weighted results for the region are more affected than in the 33% station removal scenario, as the highly populated areas are impacted by the removal of that band of sensors. The warning time wasn't significantly affected by this sensor removal, similar to the 33% station removal scenario. This analysis indicates that PLUM could work well in the Canterbury region even under reasonably significant failure scenarios. The primary reason is the very high density of sensor stations around the high population area of Christchurch, which provides sufficient redundancy to cope with even substantial failures.

However, catastrophic failures that could take down the entire communication network or remove all stations within the warning radius of Christchurch City would probably cause a failure of PLUM because no stations would be available to issue a warning based on the PLUM architecture.

#### **Limitations of the analysis:**

This analysis used a synthetic earthquake catalogue, which is developed based on the active fault database of New Zealand. However, it does not take unknown/blind faults into account. Therefore, the results could not represent consequences of unknown faults such as the Port Hills Fault which caused the highly damaging Christchurch earthquake. The Port Hills Fault is a hidden fault buried under sediments with no visible signs on the surface (Even the rupture of the Port Hills Fault never reached the surface for the Christchurch earthquake). Because this fault was so close to the city, it would have been impossible for PLUM to have provided timely warnings to the high population centre. Therefore, the results of the analysis are optimistic as the analysis did not consider any faults near or under Christchurch City. Future

work should consider synthetic catalogues that include earthquakes due to unknown and blind faults as well as mapped faults.

The PLUM confirmation strategy used in the analysis of the Canterbury region required the confirmation station to be within a certain distance from the target grid location (80 km from the grid location in the analysis). This meant that stations further away from the grid location could not provide the confirmation even if they were close to the detection station. For some earthquakes, this strategy could introduce an unnecessary delay in the alerting process.

To address that limitation, the algorithm could use an alternative confirmation strategy that requires the confirmation station to be within a certain distance from the detection station, rather than from the grid location. A preliminary analysis using this strategy indicated that there is no major difference in the performance of the two strategies for the Canterbury region, so the full analysis was not repeated using the second confirmation strategy. Chapter 5 will show that this was not the case for the Wellington region, where the first confirmation strategy was unable to provide confirmation for alerts coming from the South Island sensors since there were so few sensors in the South Island that were close enough to Wellington. Therefore, the Wellington region analysis uses the second confirmation strategy addressing that limitation in this analysis.

### **Site class assumptions**

The site classes of the grid locations used in this analysis were chosen to be the vs30 site class of the location of the centre of the grid cell. This choice was based on the pattern of the Japan EEW system that has observation stations installed at the target region (i.e., grid cells). If the VS30 site class varies considerably across the grid cell, choosing the VS30 site class of the centre of the grid cell may be unrealistic. For the Canterbury region, the VS30 site classes were reasonably uniform across the grid cells, especially near the high population centre, so the assumption is reasonable. For the Wellington region, there is greater variation in the VS30 site classes so that simply choosing the site class at the centre of the grid cell would be unreasonable. Chapter 5 therefore uses a more complex method to define the site classes for the grid cells.

Further discussion of points that need consideration of both Wellington and Canterbury regions are presented in Chapter 6.





# Chapter 5 Evaluating PLUM on the Wellington Region

## 5.1 Introduction

The Greater Wellington region was the second region selected to evaluate the PLUM algorithm. The region is a significant earthquake-prone area in New Zealand, and it covers both low and highly populated areas. After Auckland and Christchurch, the Wellington metropolitan area (including Wellington City, Lower Hutt, Porirua and Upper Hutt) is the third most populated in the country, and Wellington City is the capital of New Zealand.

This region is vulnerable to earthquakes due to its proximity to several major geological features, including the Hikurangi subduction zone, the Wellington fault, and the Wairarapa fault. The Hikurangi subduction zone is where the Pacific Plate is being forced beneath the Australian Plate. This zone extends from the east coast of the North Island to the North end of the South Island, including the Wellington region in its Southern Hikurangi Margin. Earthquakes in this subduction zone can produce some of the largest and most destructive seismic events. The Wellington fault is a major active fault running through the city of Wellington. The Wairarapa fault lies east of Wellington and caused an 8.2 magnitude earthquake in 1855. In addition, the Cook Strait region, which separates the North and South Islands of New Zealand, is also an area of significant seismic activity, which could affect the Wellington region. For instance, the 2013 M 6.5 Cook Strait earthquake occurred in this area, followed by the 2013 M 6.6 Lake Grassmere earthquake just south of the Cook Strait. The Kaikōura earthquake, which happened in 2016 in the northern part of the South Island with a magnitude of 7.8, also impacted Wellington, causing damage to the infrastructure and buildings (Cubrinovski et al., 2020).

Therefore, the Wellington region is an important region to test the PLUM EEW. As there are offshore seismic activities (both in the Southern Hikurangi Margin and Cook Strait, it will be useful to evaluate how PLUM EEW will work using only onshore sensors. Japan has both onshore and offshore sensors (Dhakal et al., 2017), yet NZ does not have an offshore sensor network that can provide live continuous seismic data.

The analysis of the Wellington region in this chapter presents several significant differences compared to the Canterbury region analysis discussed in Chapter 4. First, Wellington features more complex geology and geography, with a broader variety of site classes. This complexity makes selecting appropriate site classes for the grid cells more challenging. Additionally, Wellington has large active faults located much closer to the population centres, which could potentially reduce the warning time provided by an Earthquake Early Warning (EEW) system and increase the potential damage. The proximity of population centres to the ocean in two directions further complicates the situation by reducing the number of sensors between some major faults and the population centre, thereby making it more difficult for the PLUM method to provide an effective warning.

The PLUM analysis for the Wellington region follows a similar methodology to that used for the Canterbury region, with some modifications in how site classes are determined, and how confirmation messages are generated. These changes to the method are described in Section 5.2. The results are presented in Section 5.3, and the discussion is in Section 5.4.

## **5.2 Methodology**


### **5.2.1 Region of interest**

The region of interest was selected by considering the boundary of the base map of the official Greater Wellington Region (*Greater Wellington Regional Council, 2024*). The main reason for selecting this area was to cover the Wellington metropolitan area, as well as regional centres like Kāpiti and Masterton, along with the lower-density areas of the southern, western and eastern Wairarapa to ensure a broad coverage of both urban and rural regions. The study included stations that could either provide an alert or confirm alerts from other stations, with coverage extending up to 250 km from the centre of the region.


Figure 5-1 shows the selected region which is divided into  $0.1^\circ \times 0.1^\circ$  grid cells. The region includes a total of 109 grid cells (indicated by red dots) all of which are at least partially covered by land. Those grid cells were considered as the receiving locations of the alerts, and midpoints of the same grids were used to simulate ground motion.

It was assumed that the ground motion throughout each grid cell is similar to that at its midpoint, that every location within a grid cell receives the same alert, and that the  $0.1^\circ \times 0.1^\circ$  grid cell boundary serves as the geofencing for the alert. Geofencing refers to the use of virtual geographic boundaries, defined by coordinates, to trigger specific actions, such as sending an alert when a mobile device is in that specific area. Note that, in the actual implementation the geofencing for the alerts would be different. If the alerts were distributed via the cellular network, the geofencing would be the areas covered by each cell in the cellular network (typically have a hexagonal shape in cellular networks). If different communication technologies were used to disseminate the alert messages, the geofencing would be different again. Since the final alert broadcasting is not within the scope of this thesis, this analysis assumes the  $0.1^\circ \times 0.1^\circ$  grid cell as the geofencing boundary.

Figure 5-2 shows the locations of all the GeoNet strong motion sensor stations that could possibly issue or confirm alerts to the selected region, given a warning radius of up to 70 km and a confirmation radius of 80 km. Synthetic seismograms are generated for all the grid points in Figure 5-1 and all the sensor stations in Figure 5-2.



*Figure 5-1 The selected Wellington region with 0.1° x 0.1° grid. Selected region is the area inside the red rectangle and covers the Greater Wellington region. Grid cells containing red dots are the locations where alerts are received. Significant population centres are labeled in black. Note that some stations for issuing and confirming alerts are outside the red region; see Figure 5-2 for station locations.*



*Figure 5-2 Wellington region with sensors. Red rectangle shows region containing all the grid locations at which the PLUM algorithm is to be evaluated. Black triangles indicate all the GeoNet permanent strong motion sensors that could possibly issue or confirm alerts to the central region.*

### **5.2.2 Site classification**

The site classification of the 109 grid locations is required for the ground motion simulations. The VS30 site classification model (Perrin et al., 2015) is used to define the site classification for the grid locations. The resolution of the VS30 site classification model is  $0.01^\circ \times 0.01^\circ$  and, therefore, it is necessary to work out the site class for each of the 109 grid locations from the one hundred VS30 site classes within each grid cell. For Canterbury, the site class was generally close to uniform across each grid cell, hence taking a VS30 value next to the centre of the grid cell was appropriate. The greater complexity of the geography in Wellington means that determining site classes for the grid points is more complicated. Particularly around the densely populated Wellington metropolitan area, there is considerable variation in the site class across each grid cell and simply choosing the VS30 site classification of a point next to

the centre of the cell is not appropriate. Consequently, several different schemes were considered. One scheme computes the average of the VS30 data values across the grid cell; nevertheless, since the damage is most likely to be where the population is focused, this scheme may underestimate the shaking. Therefore, computing a population-weighted average of the VS30 data values across the grid cell was considered to be the most appropriate for evaluating the EEW system.

Figure 5-3 presents the VS30 site classification values across the Wellington region, measured in meters per second (m/s). VS30 refers to the average shear-wave velocity within the top 30 meters of the Earth's surface. Areas marked in white to yellow colours correspond to lower VS30 values, indicating softer soils with shear-wave velocities typically below 500 m/s. Those locations with lower VS30 values are more susceptible to seismic wave amplification. Locations marked in orange are composed of stiffer soils or a mix of soft soil and rock, which has moderate amplification of seismic waves. Red to dark red indicates the highest VS30 values, typically found in areas with very stiff soil or bedrock and they are least prone to seismic wave amplification.

The population of the region is shown in Figure 5-4 with a grid resolution of  $0.01^{\circ} \times 0.01^{\circ}$  (the same as the VS30 data). The map clearly illustrates a dense population concentration in and around Wellington City, extending into the western and northern suburbs of the city. Lower Hutt and Upper Hutt also have significant population densities, particularly around areas closer to the harbour. Further north, the Kāpiti Coast, covering towns such as Paraparaumu and Waikanae also show substantial population clusters. Additionally, the Wairarapa region in the east, including Masterton and Carterton, exhibits noticeable population pockets. Both the Kāpiti Coast and Wairarapa clusters are less dense compared to the urban centres closer to Wellington. The two grid cells covering most of Wellington City (the grid cell at coordinates (-41.25, 174.75) and the one below it) are the two most densely populated grid cells in the region.




Figure 5-3 VS30 site classification values across the Wellington region. Values shown for each  $0.01^\circ \times 0.01^\circ$  data cell (plotted using the model of Perrin et al., 2015).

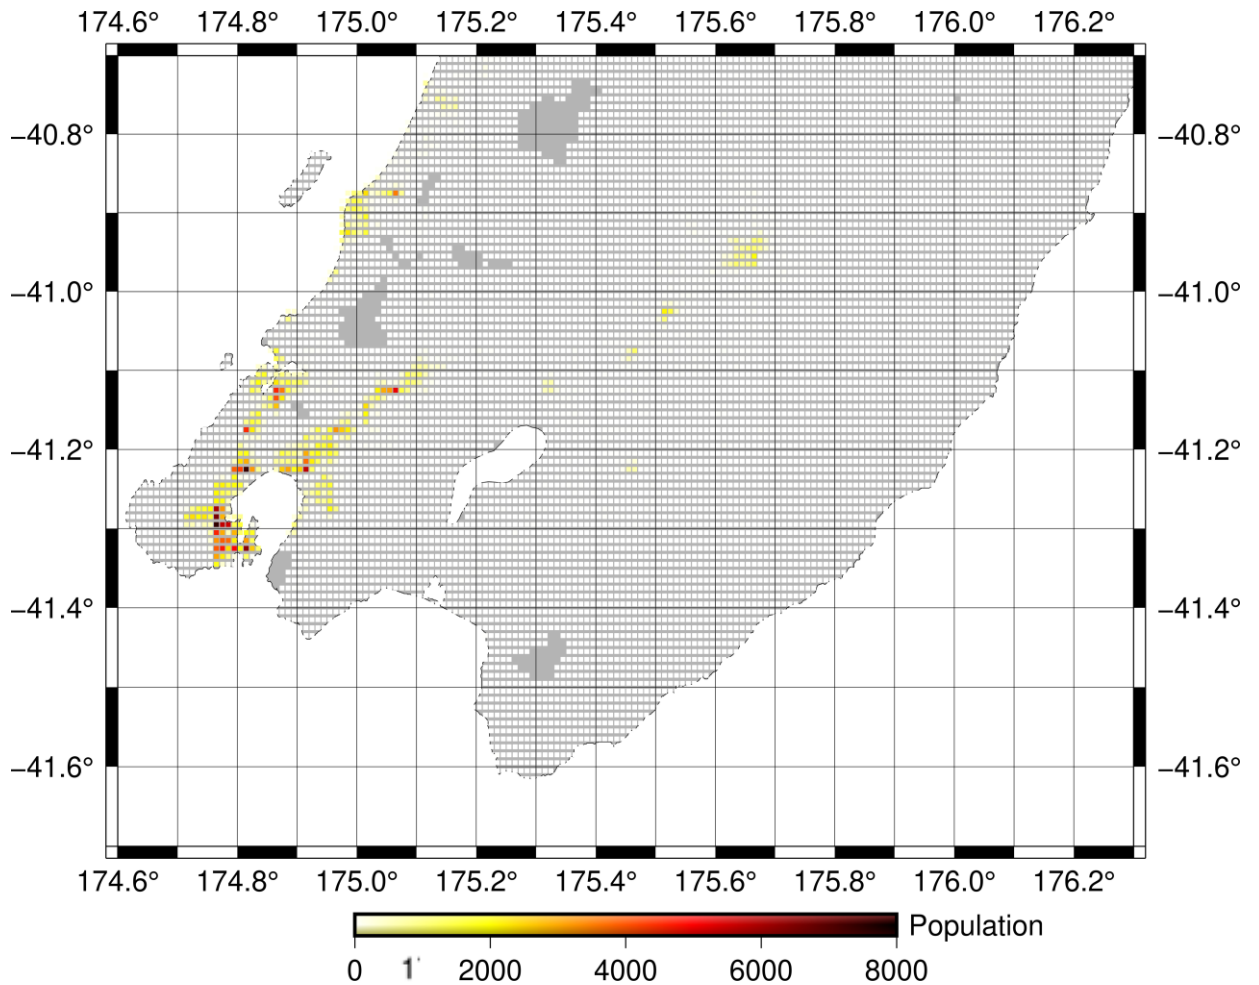


Figure 5-4 Distribution of population in Greater Wellington region. Small squares show the population in each 0.01°x0.01° data cell. Grey areas indicate zero population. The larger squares show the grid cells of the region. (Source of population data: WorldPop., 2020).

Figure 5-3 and Figure 5-4 show that there are considerable variations in population and VS30 site classes across some grid cells in the region, particularly in the grid cells with high populations in Wellington City and the Hutt Valley (Lower Hutt and Upper Hutt). Therefore, to compute an appropriate site class for each grid location, each VS30 value in a grid cell is weighted by the population of the corresponding 0.01° x0.01° data cell to determine the population-weighted VS30 value for the grid cell,  $VS30_g$  (Note that offshore points are automatically excluded because the population values are 0) as shown in Equation 5-1.

$$VS30_g = \frac{\sum_{sg}(pop_{sg} \times VS30_{sg})}{\sum_{sg}(pop_{sg})} \quad (5-1)$$

where  $g$  is a grid cell, and  $sg$  ranges over all 0.01°x0.01° data cells inside  $g$ .

From this population-weighted average  $VS30_g$  value, the site class of the grid cell is determined according to the following criteria (NEHRP, 2023), as in Section 4.2.2 of Chapter 4.

- $VS30_g > 1500$  m/s → Class A (Hard Rock):
- $1500 \geq VS30_g > 760$  m/s → Class B (Firm to Hard Rock)
- $760 \geq VS30_g > 360$  m/s → Class C (Dense Soil and Soft Rock):
- $360 \geq VS30_g \geq 180$  m/s → Class D (Stiff Soil):
- $180 > VS30_g > 0$  → Class E (Soft Soil):

Figure 5-5 illustrates the site classes of all the grid locations, computed using both the population-weighted scheme (red letters) and the unweighted scheme (blue letters). The difference between the two schemes is clear, particularly in the grid cells in Wellington City and Hutt Valley. The population-weighted site classes are used for the ground motion simulation.




Figure 5-5 Computed grid location site classes: Red = population-weighted  $VS30$  mean; Blue = unweighted  $VS30$  mean.

The station site classes were obtained from the station site classifications provided by Kaiser et al. (2017b) and Wotherspoon et al. (2024). However, for sensors installed in buildings or locations with limited access, or newly installed stations, the site class information is unavailable in either dataset. In such cases, the closest VS30 site class is used.

### 5.2.3 Selection of sub-catalogues from the main catalogue.

As in the approach used for Canterbury, one hundred 100-year sub-catalogues were obtained from the RSQSim synthetic earthquake catalogue and used to analyse the PLUM algorithm for the Wellington region (The 100-year periods were the same as the ones used for the Canterbury analysis in order to facilitate comparison between the results). Out of the 3715 earthquakes (magnitude greater than 6) that occurred in New Zealand over those 10,000 years, 1306 earthquakes which could affect the Wellington region were selected. All 1306 earthquakes are shown in Figure 5-6 and the earthquakes from one 100-year period are indicated in red colour circles.

1272 earthquakes were considered to be crustal events, and 34 to be Hikurangi subduction events. Hikurangi subduction earthquakes were identified by calculating the fault plane geometry of all M7.5+ earthquakes. It was assumed that Hikurangi faults had a fixed strike angle of 210°, and a dip angle of 10°; Wells and Coppersmith's equations (as shown in Equations 5.2 and 5.3) were used to calculate the fault length and width (these equations are the same as Equation 4.2 in Chap 4).

$$\text{Fault length (km)} = e^{(-2.42 + 0.58M)} \quad (5.2)$$

$$\text{Fault width (km)} = e^{(-1.61 + 0.41M)} \quad (5.3)$$

where  $M$  is the magnitude.

The fault's top right corner was then calculated from the centroid of the fault and the fault length and width. Any of these earthquakes whose fault's top right corner was offshore from the east coast were assumed to be Hikurangi subduction earthquakes and were simulated accordingly, under the assumption that they could all impact the Wellington region.

All other earthquakes were assumed to be crustal earthquakes. Using the GMPE and the same criteria as specified in Section 4.2.3, 1272 crustal earthquakes that could affect the region were selected.

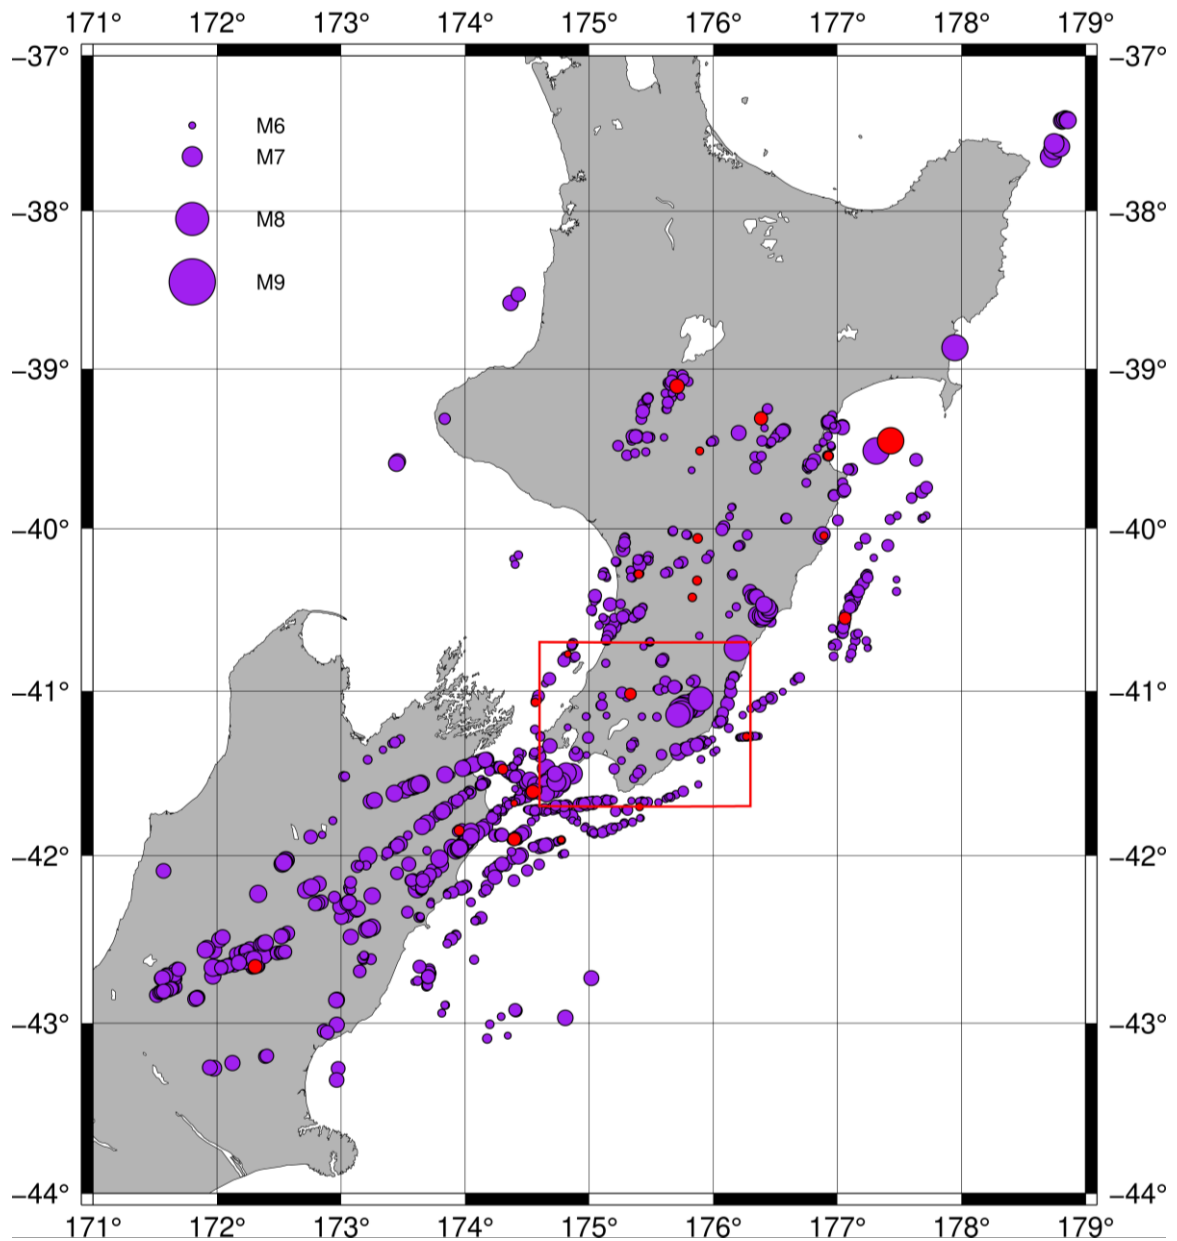


Figure 5-6 All 1306 earthquakes that could affect the Wellington region. The circles are the centroids (synthetic catalogue) of the earthquakes in the random 100 sets of the 100-year periods from the 220000-year catalogue, and the size of the circles represents the magnitude. The red circles are earthquakes from one selected 100-year period (#37). The red rectangle is the inner region containing the grid points where the performance of the EEW algorithm is evaluated.




Figure 5-7 Histogram showing distribution of number of earthquakes: each bar shows the number of periods that had a given number of earthquakes.

#### 5.2.4 Synthetic seismogram generation for the catalogue events

The subduction and crustal earthquakes were simulated using EXSIM to generate synthetic seismograms following the same methodology described in Section 4.2.4 of Chapter 4. Most of the earthquake simulation parameters are the same as in the Canterbury experiments, however, the Q-model, which represents the attenuation of seismic waves, uses parameters specified for the Wellington region (Holden et al., 2013; Carey et al., 2022).

#### 5.2.5 PLUM algorithm

Initial simulation in the Wellington region showed that the PLUM algorithm had difficulty providing correct timely alerts to Wellington City because of the distance from the sensors towards the south and west, which were located on the South Island. The initial theory was that this might be caused by the confirmation strategy used in the PLUM experiments for the Canterbury region, where the confirmation for an alert signal was received by the grid point from a second station within the confirmation radius of the grid point. Many of the South Island sensor stations were too far from Wellington to provide confirmation. The confirmation strategy was therefore modified to require the alerting station to receive a confirmation signal from a sensor station within the confirmation radius of the alerting station, rather than the grid point. This strategy also provides confirmation that an alert is not due to noise or error

at the alerting station, but uses a different set of stations to provide the confirmation. Although the description of the confirmation strategy for the Southern California region (Cochran et al., 2019; Kilb et al., 2021) is a little ambiguous, the revised strategy seems to be more in line with the Southern California strategy.

The PLUM algorithm with confirmation from the confirming station to the alerting station is as follows (refer to Figure 5-8).

### PLUM algorithm with confirmation to alerting station

- *When a station detects shaking (PGA at  $\tau_l$  or above), it issues a confirmation message to all stations within the confirmation radius.*
- *When a station detects strong shaking (PGA at  $\tau_h$  or above), then*
  - *If the station has already received a confirmation message (within the confirmation window), it will immediately issue an alert to all points within the warning radius.*
  - *Else, it waits, and if it receives a confirmation message before the end of the confirmation window, it will then issue an alert to all points within the warning radius.*
  - *Else, if it does not receive a confirmation message within the confirmation window, it will abandon the alert.*




Figure 5-8 PLUM algorithm with confirmation from confirming station to alerting station.

Just as in Chapter 4, it is simpler (and more efficient) for the simulation of the PLUM algorithm to “invert” the algorithm above, focusing on the grid points, rather than on the sensor stations that issue the alerts. The algorithm used for this simulation of PLUM was as follows:

## PLUM Algorithm with confirmation to alerting station, for analysis

### Parameters:

- $\tau_h$  (strong shaking detection threshold)
- $\tau_l$  (weak shaking detection threshold)
- warning radius
- confirmation radius
- comm-time [communication latency]

### Output:

- outcome and warning-time for each grid point for each earthquake.

### Algorithm:

**For each** grid point,

    compute list of stations within warning radius of grid point

**For each** station

    compute list of nearby stations within confirmation radius of station

**For each** earthquake

**For each** station

        ws-time[station] = first time when  $pga \geq \tau_l$  (infinity if no such time)

        ss-time[station] = first time when  $pga \geq \tau_h$  (infinity if no such time)

[Determine time confirmation or alert is issued (infinity if no alert issued)]

    Initialise conf-time = infinity, alert-time[station] = infinity.

**for each** other station in list of stations nearby to station

        conf-time = min(conf-time, ws-time[other station])

    alert-time[station] = max(ss-time[station], conf-time)

**For each** grid point

    ws-time[grid-pt] = first time when  $pga \geq \tau_l$  (infinity if no such time)

    ss-time[grid-pt] = first time when  $pga \geq \tau_h$  (infinity if no such time)

    alert-time = infinity

**For each** station on list of stations within warning radius of grid point

        alert-time = min(alert-time, alert-time[station])

[Determine outcome of event at grid point:]

**If** ss-time[grid point] < infinity **then** [ie, strong shaking at grid point]

**If** alert-time = infinity **then**

            outcome = MA

**Else If** alert-time + 2s < ss-time[grid point] **then**

            outcome = CTA

            warning-time = ss-time – alert-time – 2s

```

Else
    outcome = LA
    warning-time = ss-time– alert-time-2s
Else If ws-time [grid point] < infinity then [only weak shaking at grid]
    If alert-time < infinity then
        outcome = PA
Else [ie no noticeable shaking at grid point]
    If alert-time < infinity then
        outcome = FA
Else
    outcome = none

```

Very few of the stations on the upper South Island closest to Wellington City were within the 60 km maximum warning distance used in the Canterbury experiment. This could limit the possible alerts for Wellington City, and therefore, the experiments for the Wellington region expanded the range of warning radii up to 70 km. Note that the “undamped” assumption underlying the PLUM algorithm is questionable at larger warning radius, so that extending the warning radius further would require a modification to the PLUM algorithm so that it no longer depended on the assumption of PLUM undamped propagation within the warning radius.

### 5.2.6 Performance Indicators

The performance indicators for the Wellington region, including the population-weighted measures, are the same as those used for the Canterbury region, as described in Sections 4.2.6 and 4.2.7.

## 5.3 Results

### 5.3.1 Results for the Wellington region by location.

The PLUM algorithm described in Section 5.2.5 was evaluated on each 100-year period, applying it to each selected earthquake of that period, and for a range of different warning radii. Because of the distance between Wellington City and the stations at the North end of the South Island, we explored a wider range of warning radii than for the Canterbury region:

30 km, 40 km, 50 km, 60 km, 70 km, 80 km, 90 km, 100 km and 150 km. The results were evaluated using the same performance measures as for the Canterbury analysis. However, the PLUM assumption of undamped propagation is clearly violated at the higher radii and the key  $CTA_{norm,g}$  performance indicator showed no change for radii above 70 km. Therefore, the analysis below only considers the results for warning radii from 30 km to 70 km.

### **Grid Cell measures: $CTA_{norm}$ , $FA_{norm}$ , and $PA_{norm}$**

Figure 5-9 shows the expected  $CTA_{norm,g}$  performance, based on all 100 periods, using a PLUM radius of 30 km, 40 km, 50 km, 60 km, and 70 km. As in Chapter 4, the size of the squares represents the expected number of shaking events above the  $\tau_h$  threshold for the grid cell and the colour of the squares indicates the fraction of correct timely alerts ( $CTA_{norm,g}$ ) for those strong shaking events. Dark blue means 100%  $CTA_{norm,g}$  —the given grid point received correct timely alerts for 100% of the shaking events. Dark red means 0% CTA. Almost all grid cells of the selected region experienced some shaking above the threshold in at least some of the periods. However, most of the shaking (represented by larger squares) occurs in the coastal region because most of the site classes of those locations are defined as site class E (see Figure 5-5). The south and southeast regions, including Wellington City and Wairarapa, experienced more missed or late alerts as indicated by pink and red colour squares. The northwest region, including Porirua, Kāpiti, Lower Hutt, and Upper Hutt received more correct timely alerts as  $CTA_{norm}$  is higher than 50% for those locations, especially for the larger warning radii. The upper northeast corner of the region also received more CTA. In the Wellington City area, some locations received more CTA and some fewer. Especially, the southern part of Wellington City always receives less than 50%  $CTA_{norm}$ . The  $CTA_{norm}$  ratio is improved by increasing the warning radius. According to Figure 5-9, 70 km will be the best warning radius to obtain a high  $CTA_{norm}$  ratio. However, the population needs to be taken into account as some locations have high population and others have less, so the impact will be different according to the population. Section 5.3.2 presents the results taking the population of each grid cell into consideration.

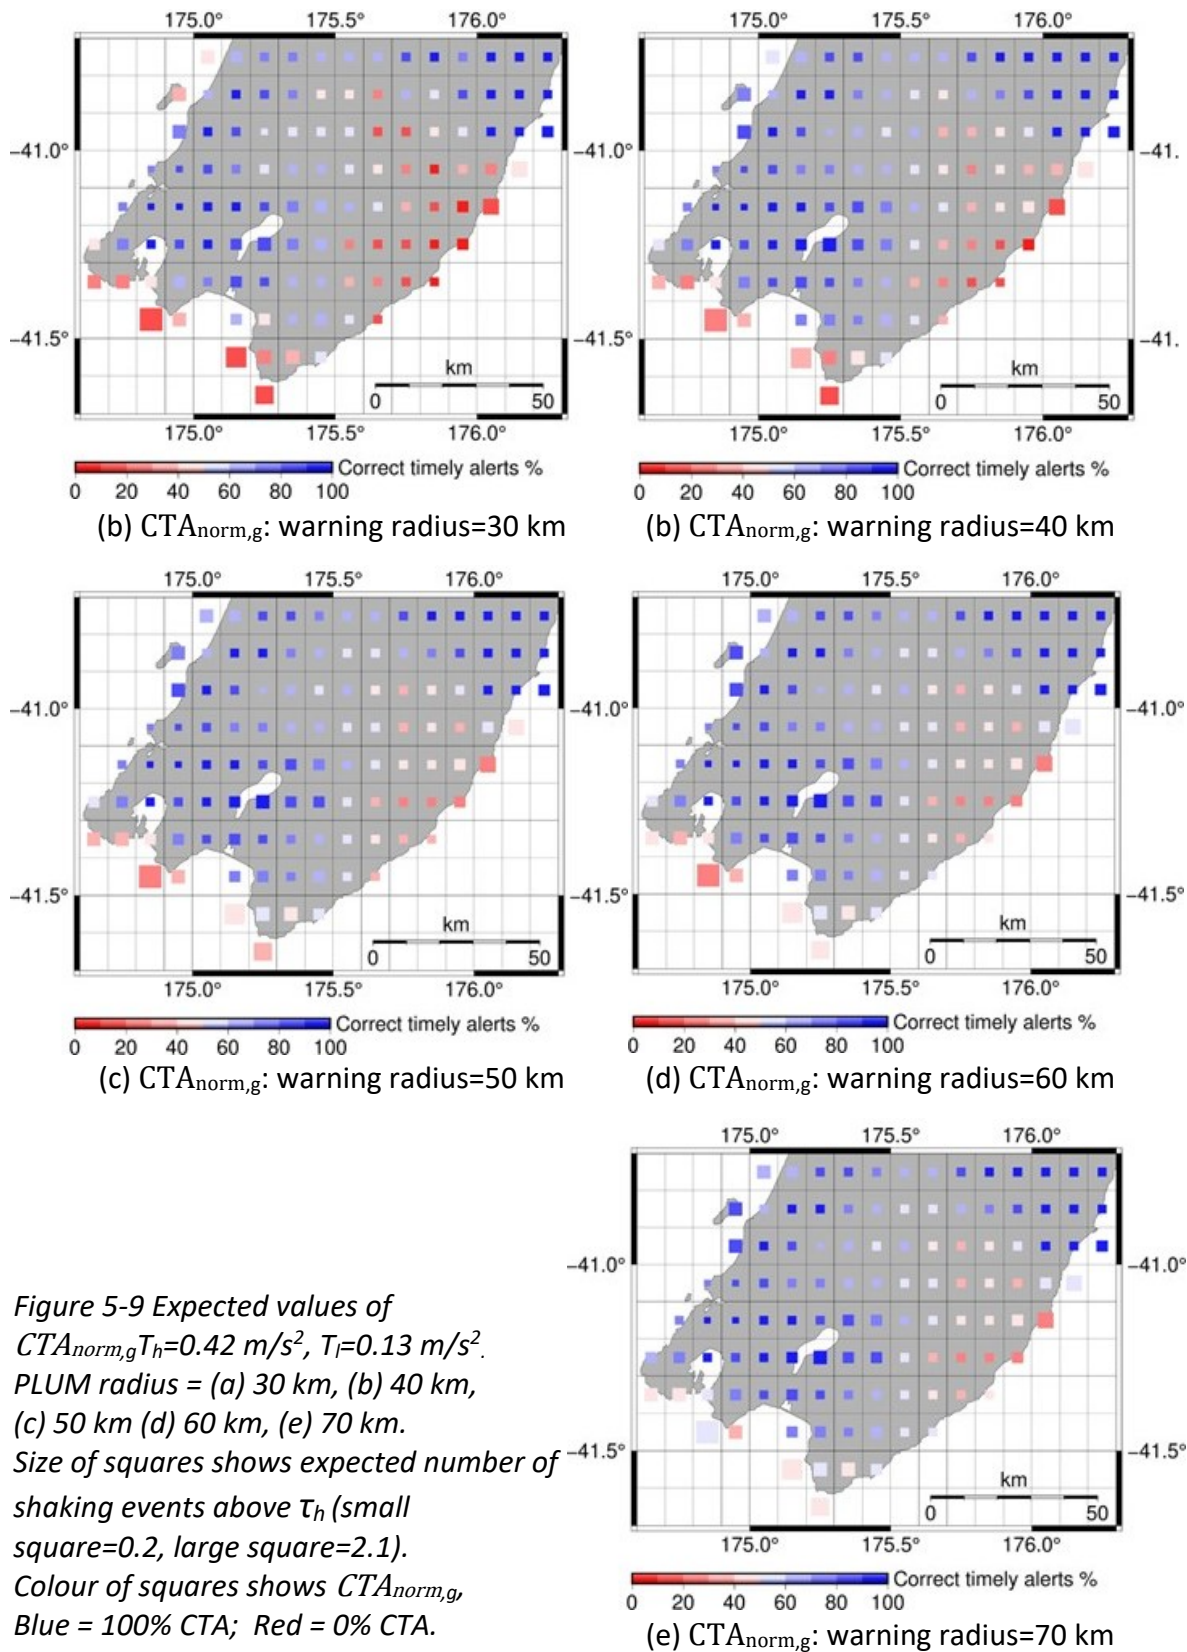
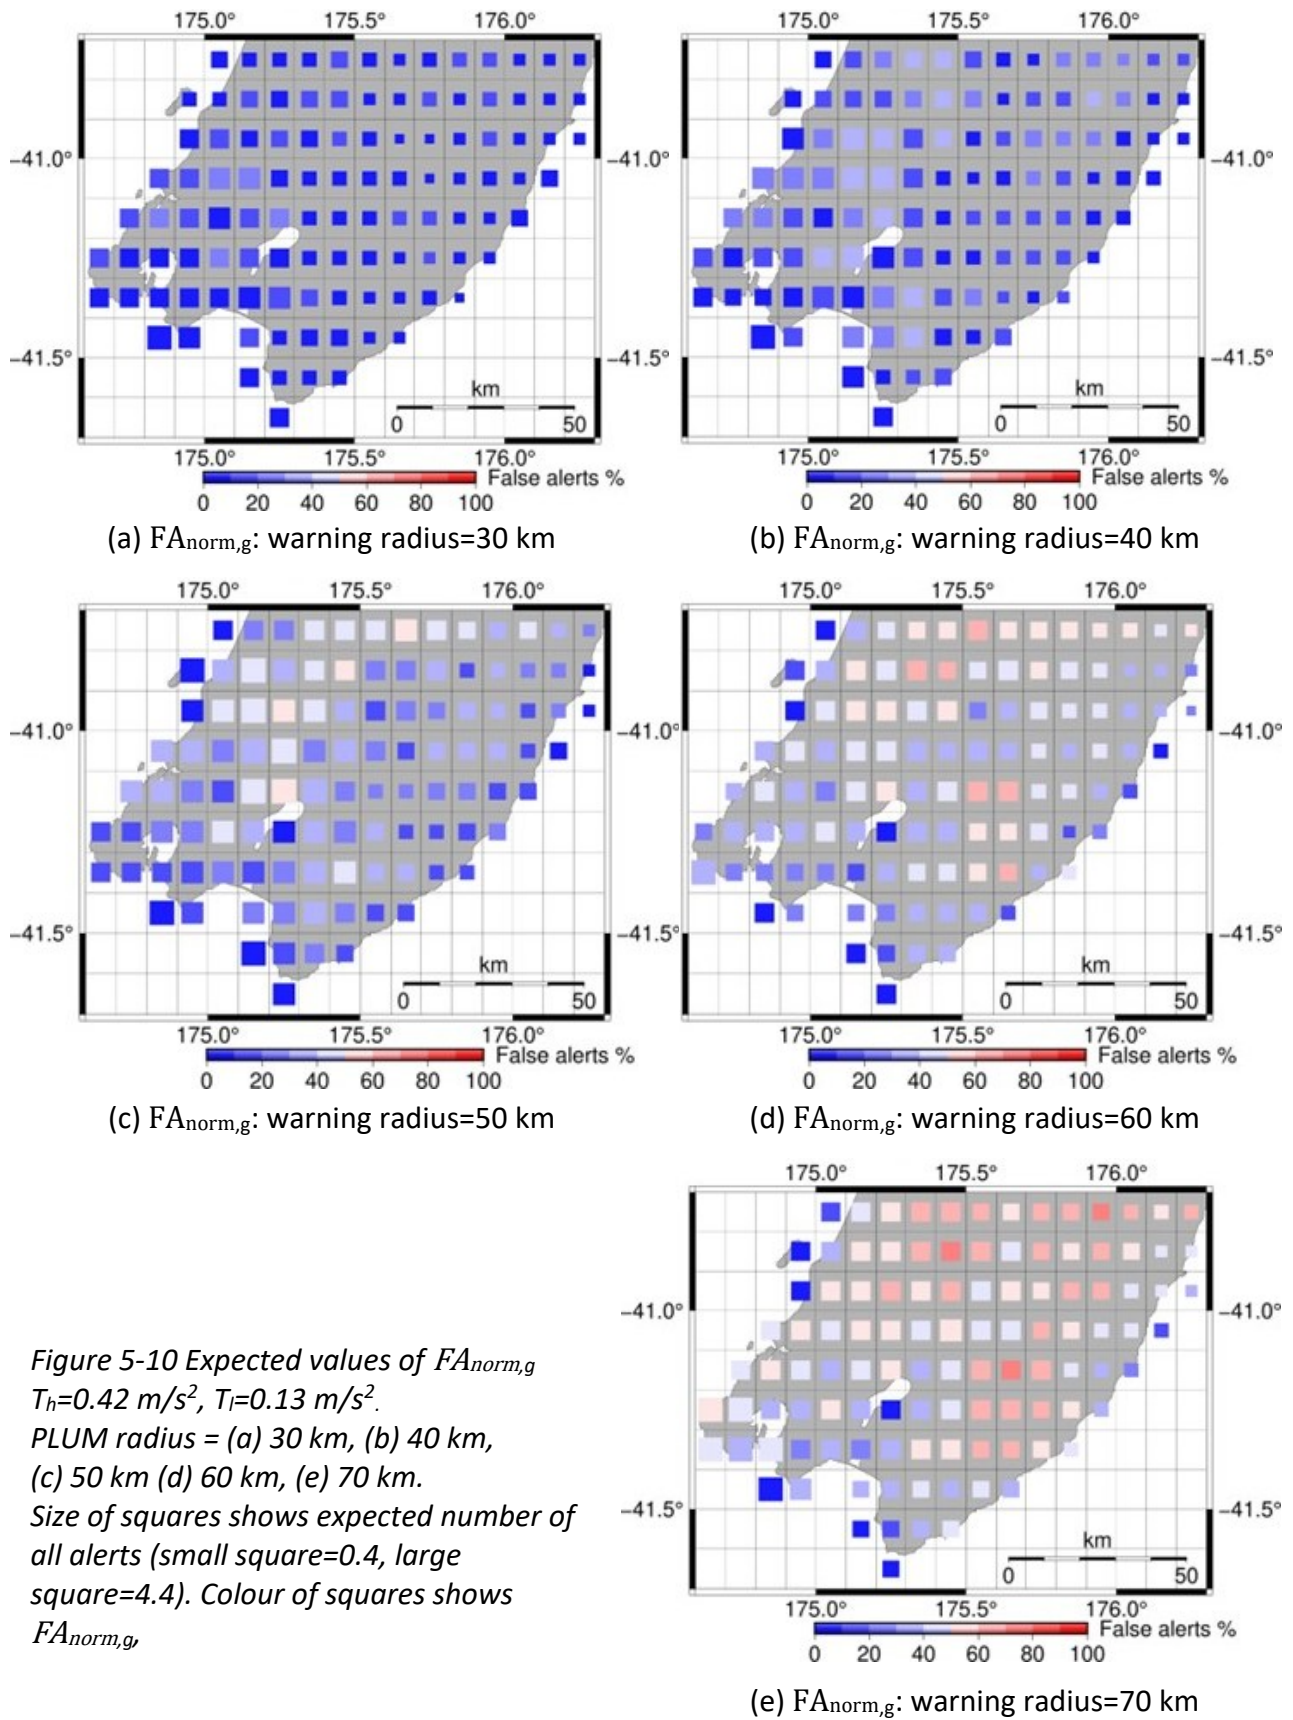
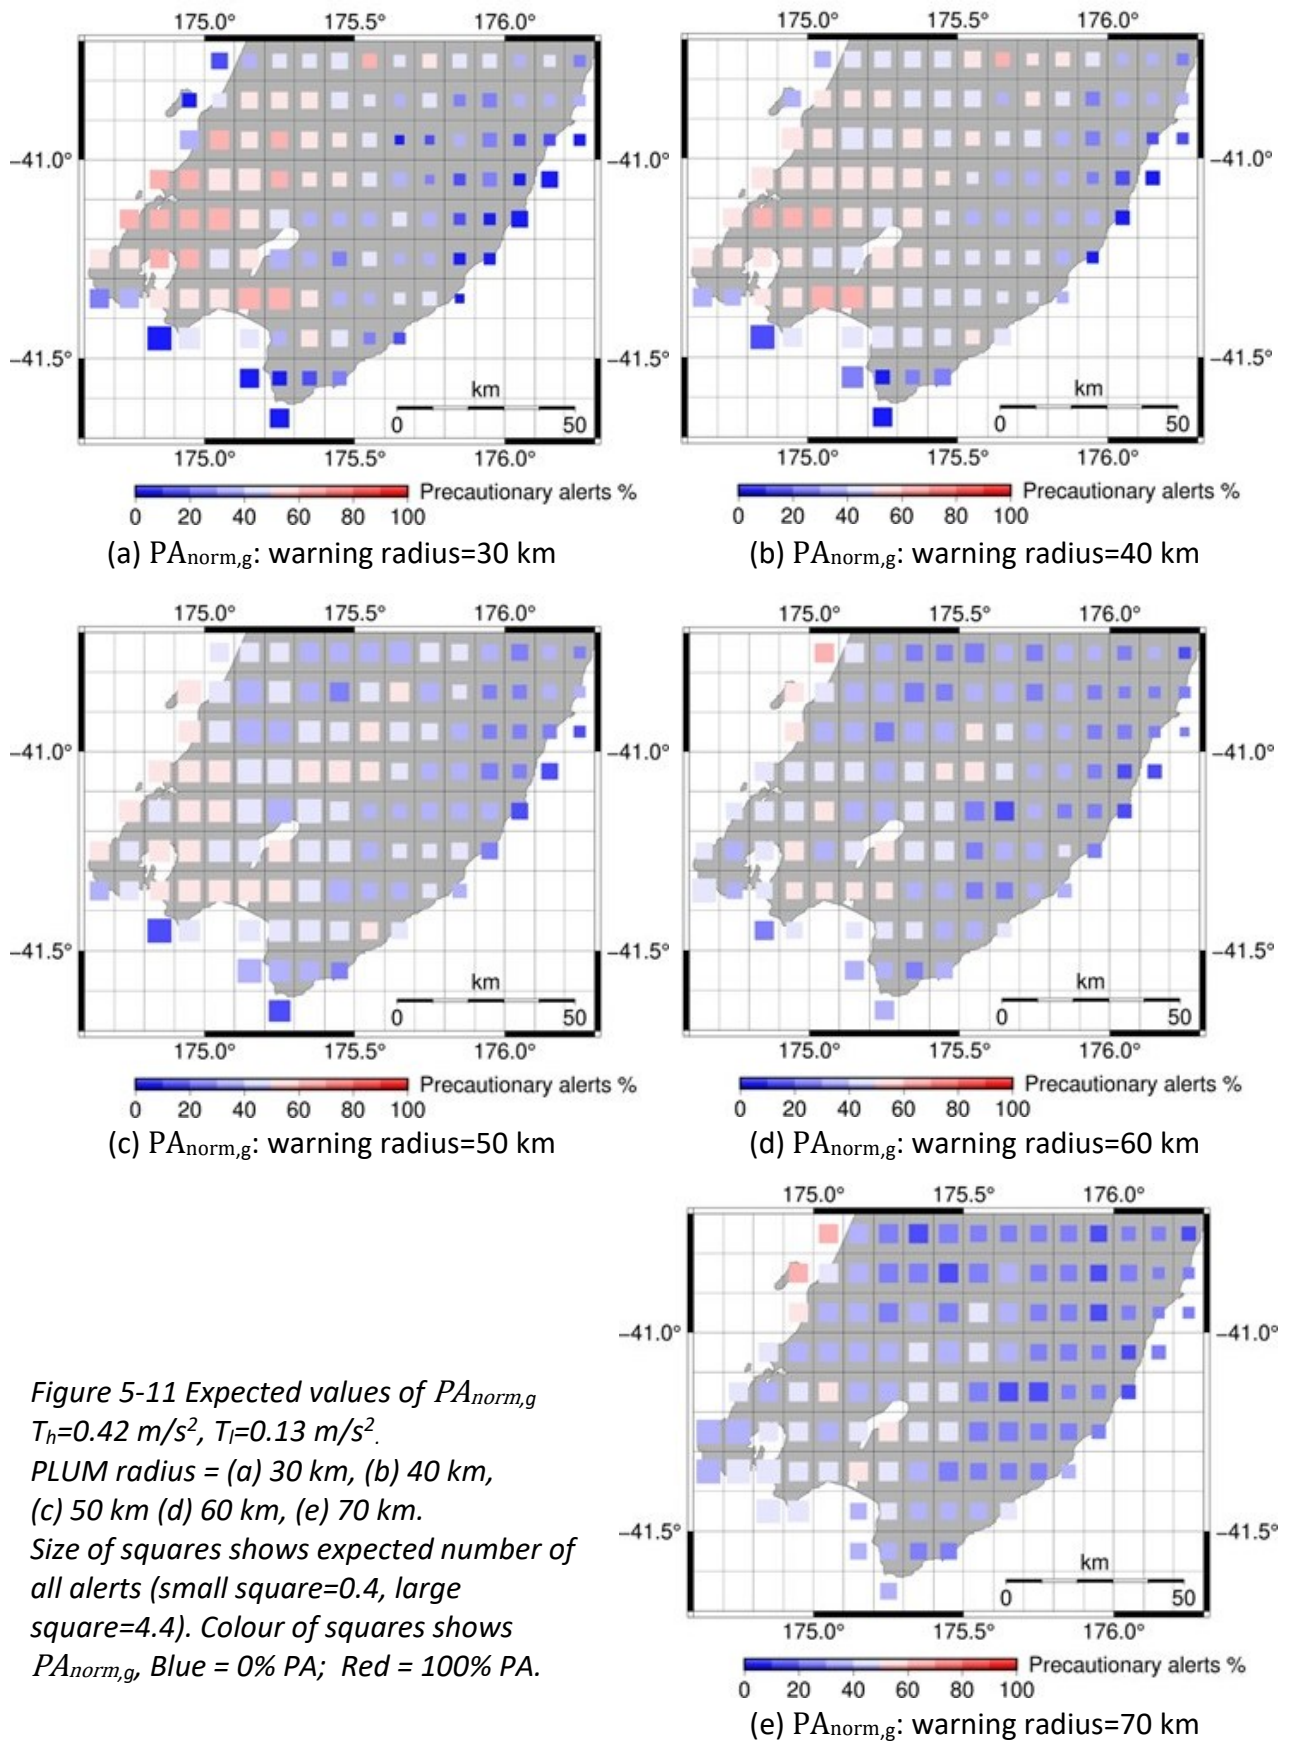


Figure 5-9 Expected values of  $CTA_{norm,g}$   $T_h=0.42 \text{ m/s}^2$ ,  $T_l=0.13 \text{ m/s}^2$ . PLUM radius = (a) 30 km, (b) 40 km, (c) 50 km (d) 60 km, (e) 70 km. Size of squares shows expected number of shaking events above  $T_h$  (small square=0.2, large square=2.1). Colour of squares shows  $CTA_{norm,g}$ , Blue = 100% CTA; Red = 0% CTA.

Figure 5-10 shows the values of the expected  $FA_{norm,g}$  (the fraction of alerts that are false alerts) in each grid cell. The size of the squares represents the expected number of all alerts (CTA+MLA+PA+FA) received by each grid cell. For the 30 km and 40 km warning radii, the expected  $FA_{norm,g}$  is less than 50% for all the grid cells. However, when the radius increases to 50, 60 and 70 km the fraction of false alerts increases and for some locations, more than 50% of the alerts are false alerts. This contrasts with the Canterbury region where  $FA_{norm,g}$  was below 50% for all locations and all warning radii.

Figure 5-11 illustrates the percentage of precautionary alerts over all the alerts for the region for different warning radii. There is a high fraction of precautionary alerts in the Western part of the region for the low warning radii, but the  $PA_{norm}$  declines as the warning radius increases. However, comparison with Figure 5-9 and Figure 5-10 shows that the decrease in precautionary alerts is largely due to the increase in false alerts.





## Overall region measures

Table 5-1 shows the expected counts of the different kinds of alerts across the whole region for a 100-year period and how they change with the warning radius. For the 30 km radius, the number of CTA and MLA alerts are very similar. Increasing the warning radius increases the CTA count slightly and decreases the MLA counts by the same amount. The total number of correct alerts (CTA+LA) increases slightly. The counts of both PA and FA increase with the warning radius, though the counts of FA increase much faster, consistent with the data shown in Figure 5-10 and Figure 5-11.

*Table 5-1 Expected number of each PLUM outcome across whole region for different warning radii.*

Warning radius	CTA (n)	CTA (%)	LA (n)	LA (%)	MA (n)	MA (%)	MLA (n)	MLA (%)	PA (n)	PA (%)	FA (n)	FA (%)	Total
30 km	27	21%	18	14%	7	5%	25	19%	45	34%	9	7%	131
40 km	30	19%	17	11%	5	3%	22	14%	62	39%	24	15%	160
50 km	32	16%	16	8%	4	2%	20	10%	75	38%	52	26%	199
60 km	32	13%	16	7%	3	1%	19	8%	83	35%	87	36%	240
70 km	33	11%	16	5%	3	1%	19	6%	92	31%	130	44%	293

Table 5-2 shows the expected  $CTA_{norm}$ ,  $FA_{norm}$  and  $PA_{norm}$  across the entire region. The expected  $CTA_{norm}$  improves with increasing warning radius. However,  $FA_{norm}$  gets much worse with increasing warning radius.  $PA_{norm}$  increases from 30 km to 40 km then decreases at higher radii. This is consistent with the grid-cell based measures shown in Figure 5-9, Figure 5-10 and Figure 5-11.

Compared to the Canterbury results (Table 4-6), the expected  $CTA_{norm}$  in the Wellington region is a little worse, but the expected  $FA_{norm}$  is about 8 times worse.

Table 5-2 Expected  $CTA_{norm}$ ,  $FA_{norm}$  and  $PA_{norm}$  across the whole region for different warning radii.

Warning radius	Expected $CTA_{norm}$	Expected $FA_{norm}$	Expected $PA_{norm}$
30 km	51.8%	8.8%	42.4%
40 km	58.3%	17.6%	44.9%
50 km	61.4%	28.9%	42.2%
60 km	62.5%	39.3%	37.5%
70 km	63.9%	47.5%	33.6%

### Warning time measures

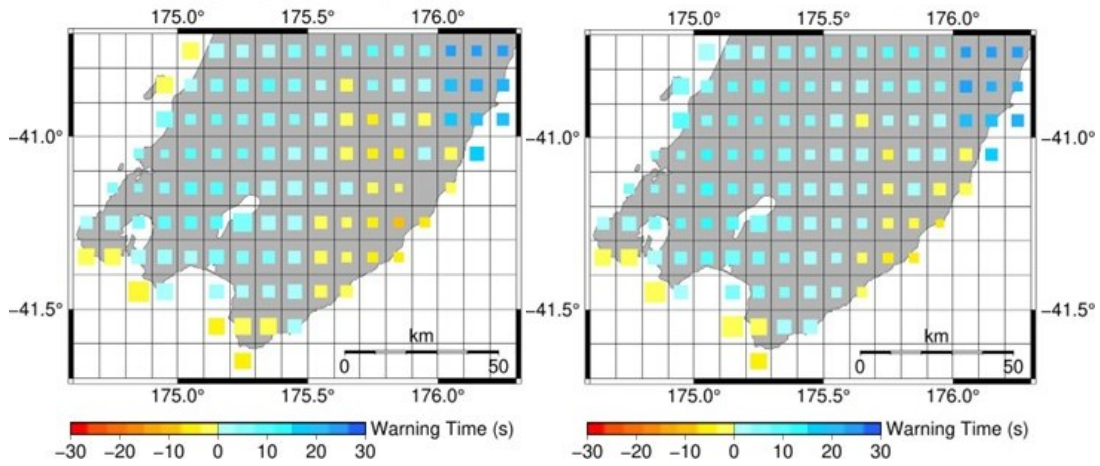
The expected warning time for each grid cell is illustrated in Figure 5-12. For warning radius of 30 km and 40 km, some locations have negative expected warning times (corresponding to many late alerts), especially in the eastern half of the region (Wairarapa) and the south coast, including the southern part of Wellington City. The largest warning times are northeast of Wellington City (including Lower and Upper Hutt) and in the far northeast corner of the region.

For the 30 km scenario, the minimum warning time is  $-8.3$  seconds (i.e., 8.3 seconds late) and the maximum is 24.5 seconds. For the 70 km scenario, almost all locations have a positive expected warning time. The minimum warning time also improves to  $-1.2$  seconds and maximum warning time increases to 29.5 seconds.

Figure 5-13 shows box plots of the distributions of the warning times of all correct timely alerts for the different warning radii. The median warning time is between 5 and 8 seconds for all cases. The warning time improves as the radius is increased from 30 km to 60 km, but there is no significant improvement at 70 km.

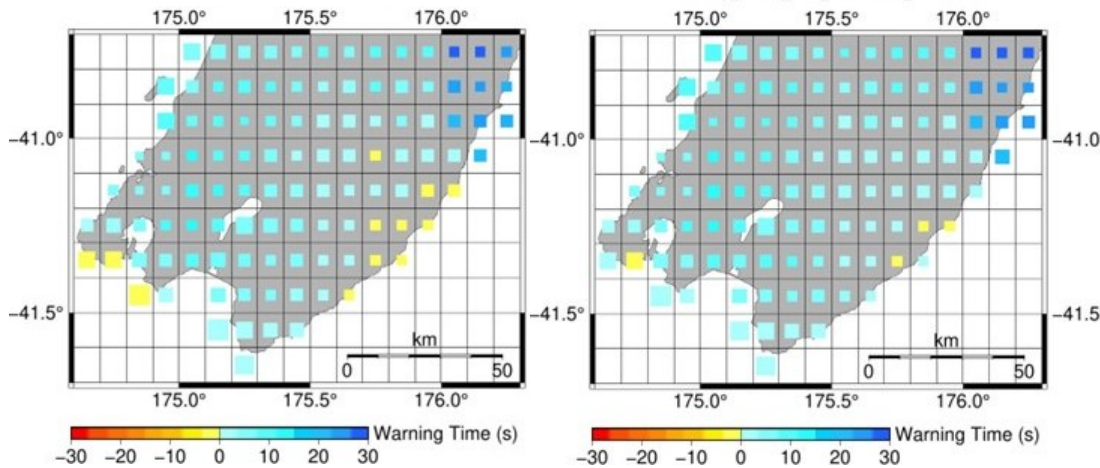
Figure 5-14 illustrates the distribution of warning times of all correct alerts (CTA and LA) for the different warning radii. The median warning time is 1.7 seconds for the 30 km warning radius and improves to 3.5 seconds for the 70 km radius.

Both sets of distributions have a long upper tail of long warning times. Table 5-3 presents the mean warning times including standard deviations. Unsurprisingly, the means are higher than the medians shown in Figure 5-13 and Figure 5-14. For example, for the 50 km scenario, the median of the correct timely alerts is 7.2 seconds, while the mean is 10.7 seconds.



(a)  $WT_{CA,g}$ : radius=30 km  
min: -8.3s max: 24.5s


(b)  $WT_{CA,g}$ : radius=40 km  
min: -7.1s max: 25.6s



(c)  $WT_{CA,g}$ : radius=50 km  
min: -2.3s max: 27.4s

(d)  $WT_{CA,g}$ : radius=60 km  
min: -1.5s max: 29.5s

Figure 5-12 Expected mean warning times for correct alerts ( $WT_{CA}$ ), for different warning radii (30 km, 40 km, 50 km, 60 km). Size of square represents the expected number of correct alerts in the cell (small square=0.19, largest square=1.81). Colour of square represents the expected mean warning time: Yellow for a negative mean warning time (late); Blue for a positive mean warning time (timely).



(e)  $WT_{CA,g}$ : radius=70 km  
min: -1.2s max: 29.5s




Figure 5-13 Distribution of *WTCTA*: box plots of warning times of all correct **timely** alerts across region. Warning radii 30 km, 40 km, 50 km, 60 km. The box plots show the median (red line), 25<sup>th</sup> and 75<sup>th</sup> percentiles, and the outliers. The dotted lines show the 25<sup>th</sup> percentile, median, and 75<sup>th</sup> percentile of the 50 km warning radius for comparison with other warning radii.




Figure 5-14 Distribution of *WTCA*: box plots of warning times of all correct alerts across region. Warning radii 30 km, 40 km, 50 km, 60 km. The box plots show the median (red line), 25<sup>th</sup> and 75<sup>th</sup> percentiles, and the outliers. The dotted lines show the 25<sup>th</sup> percentile, median, and 75<sup>th</sup> percentile of the 50 km warning radius for comparison with other warning radii.

Table 5-3 Mean  $\pm$  standard deviation of distributions of *WTCTA* and *WTCA* (warning time for timely alerts and for all correct alerts) for different warning radii.

Warning Radius	WTCTA (s)	WTCA (s)
30 km	8.8 $\pm$ 9.0	3.5 $\pm$ 9.7
40 km	9.8 $\pm$ 9.7	4.8 $\pm$ 10.4
50 km	10.7 $\pm$ 10.3	5.8 $\pm$ 11.0
60 km	11.6 $\pm$ 11.1	6.5 $\pm$ 11.8
70 km	12.0 $\pm$ 11.6	6.9 $\pm$ 12.2

### 5.3.2 Results for the Wellington region by population.

The results presented in Section 5.3.1 do not consider the population of each location. However, population is a critical parameter because it is much more important to get high CTA and low MLA, FA, and PA for the densely populated locations. Therefore, this section presents results considering the population of each grid location.

#### Population-weighted grid cell measures

Population-weighted  $CTA_{norm,g}$  is shown in Figure 5-15. The colour of the squares represents the value of  $CTA_{norm,g}$ , as in Figure 5-9, but the size of the squares is determined by the population multiplied by the expected number of strong shaking events, i.e., the expected number of person-events in the grid cell. Therefore, the diagram shows critical locations by the large squares, and the small squares (or circles when the area of the square would be less than 1% of the largest square) represent the cells with a lower number of person-events. Among the critical locations there are three grid locations (in Wellington City and Masterton) that expect correct timely alerts for less than 50% of their strong shaking events. Although the values of  $CTA_{norm,g}$  improve with increased warning radius, the most critical location (south Wellington City) still has less than 50% correct timely alerts with the 70 km warning radius. Large parts of the east of the region (Wairarapa) have poor performance, but the population is so low that this is not significant.

Figure 5-16 and Figure 5-17 show the population-weighted  $FA_{norm}$  (fraction of alerts that are False Alerts) and the population-weighted  $PA_{norm}$  (fraction of alerts that are Precautionary Alerts). The size of the squares shows the expected total number of person-alerts—(the product of the population of the grid cell and the expected number of all alerts

(CTA+MLA+FA+PA) at that grid cell. It is clear from the figures that the fraction of False Alerts around Wellington City and its suburbs, including Porirua, Lower Hutt is quite high, especially for warning radii above 50 km. This is a contrast to the Canterbury results in Section 4.3.2, where the false alert rate was very low for the Christchurch metropolitan area. On the other hand, the precautionary alert rate around Wellington City is much lower (at most 70%) than around Christchurch City because many of the incorrect alerts are false alerts rather than merely precautionary alerts.

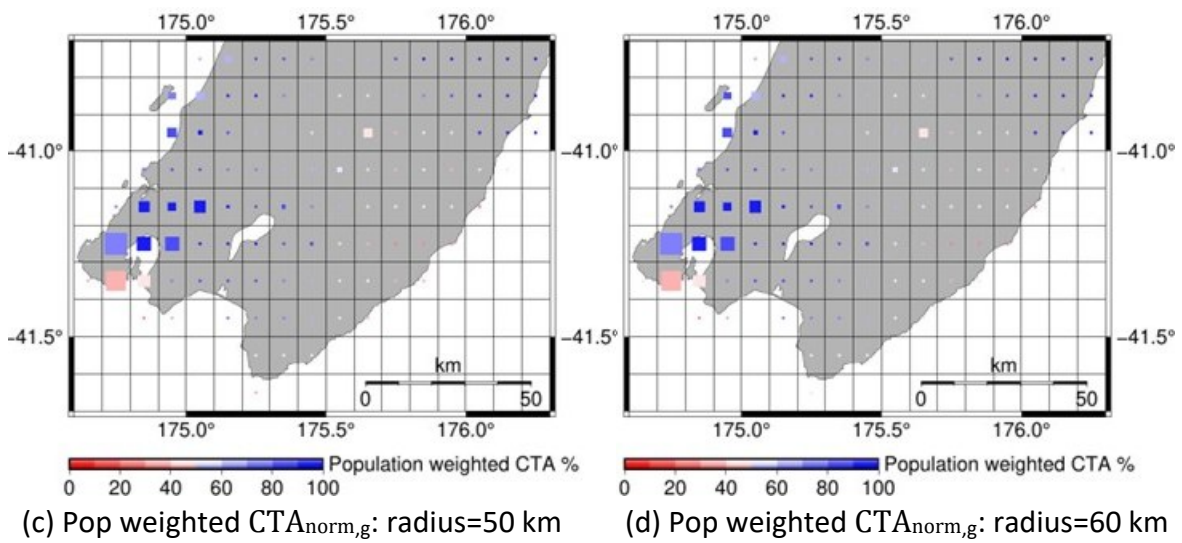
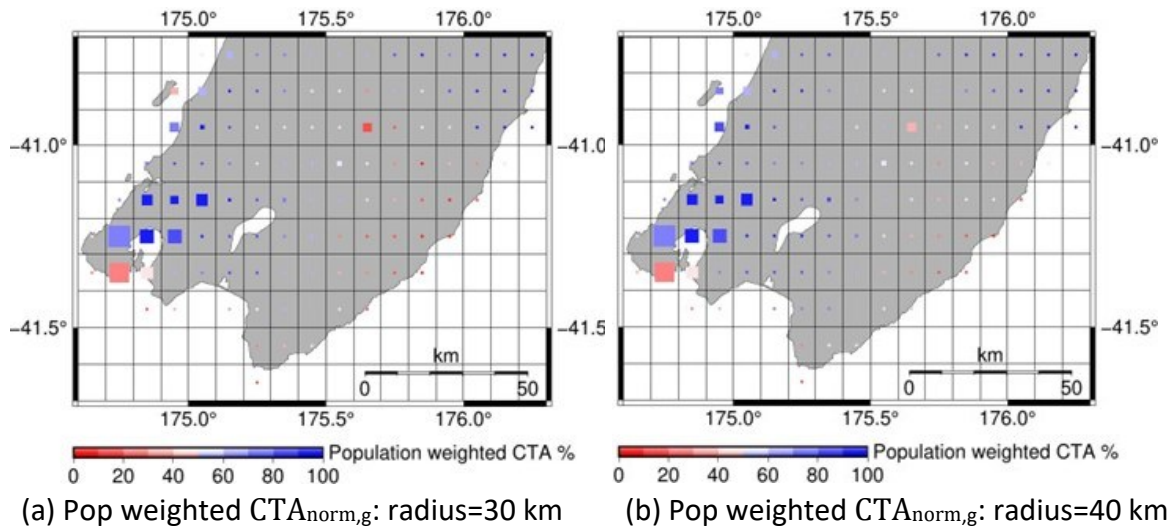

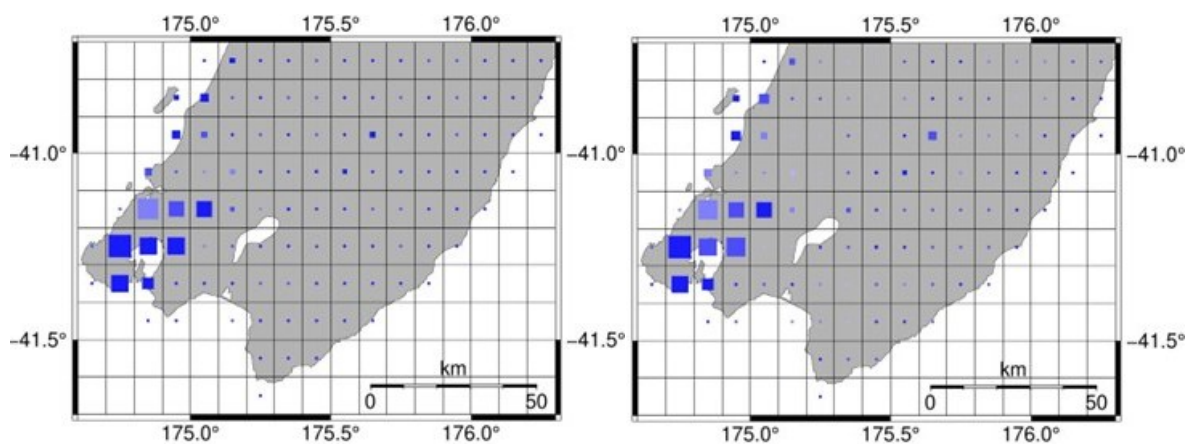


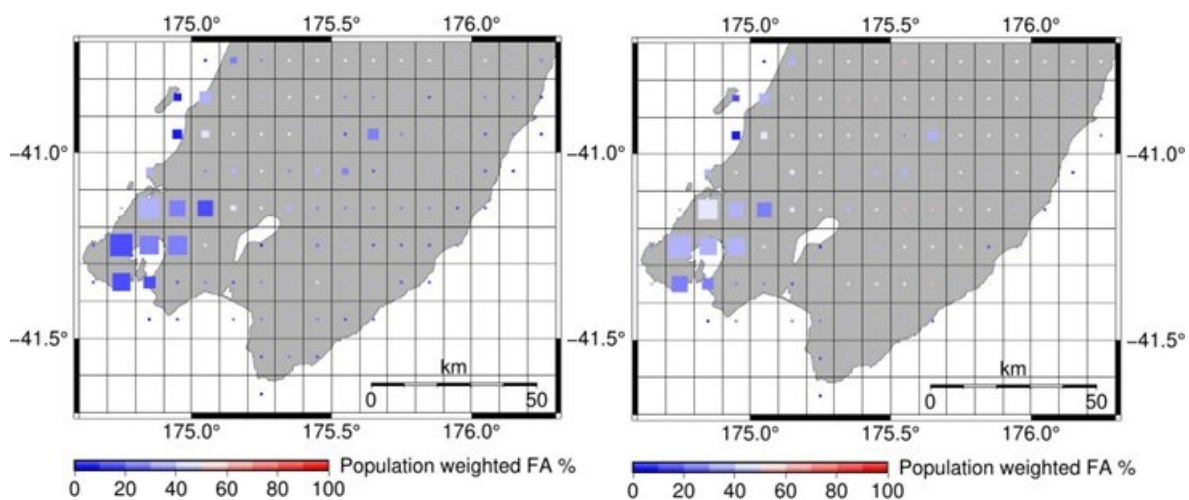
Figure 5-15 Population-weighted expected values of  $CTA_{norm,g}$   
 $T_h=0.42 \text{ m/s}^2$ ,  $T_l=0.13 \text{ m/s}^2$ .  
 PLUM radius = (a) 30 km, (b) 40 km, (c) 50 km (d) 60 km, (e) 70 km.  
 Size of square shows expected number of person-events: population  $\times$  shaking events above  $\tau_h$  (smallest square=1000, largest square=49841).  
 Colour of squares shows  $CTA_{norm,g}$ , Blue = 100% CTA; Red = 0% CTA.





(a) Pop weighted  $FA_{norm,g}$ ; radius=30 km


(b) Pop weighted  $FA_{norm,g}$ ; radius=40 km



(c) Pop weighted  $FA_{norm,g}$ ; radius=50 km

(d) Pop weighted  $FA_{norm,g}$ ; radius=60 km

Figure 5-16 Population-weighted expected values of  $FA_{norm,g}$   
 $T_h=0.42 \text{ m/s}^2$ ,  $T_l=0.13 \text{ m/s}^2$ .  
 PLUM radius = (a) 30 km, (b) 40 km, (c) 50 km (d) 60 km, (e) 70 km.  
 Size of square shows expected number of person-alerts: population  $\times$  expected all alerts at grid location (smallest square=1185, largest square=116296).  
 Colour of squares shows  $FA_{norm,g}$ , Blue = 0% FA; Red = 100% CTA.



(e) Pop weighted  $FA_{norm,g}$ ; radius=70 km

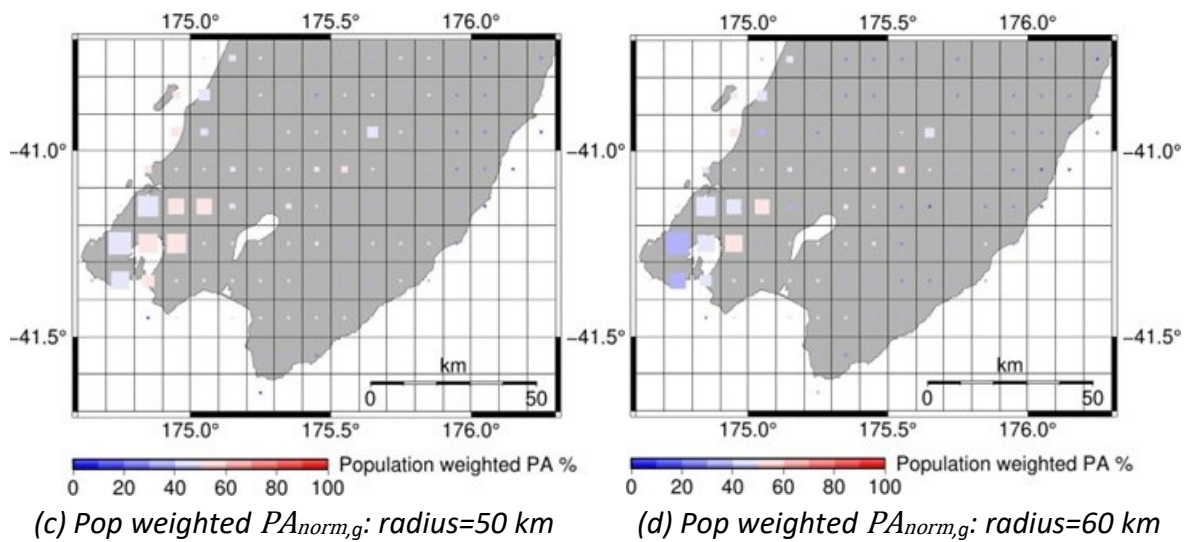
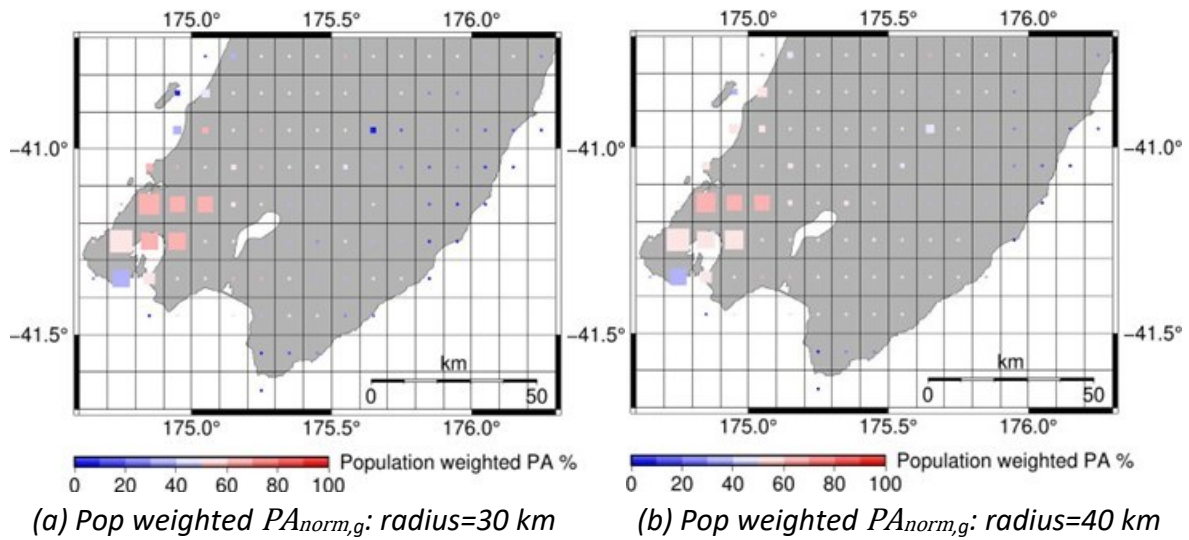



Figure 5-17 Population-weighted expected values of  $PA_{norm,g}$   
 $T_h=0.42 \text{ m/s}^2$ ,  $T_l=0.13 \text{ m/s}^2$ .  
 PLUM radius = (a) 30 km, (b) 40 km, (c) 50 km (d) 60 km, (e) 70 km.  
 Size of squares shows expected number of person-alerts: population  $\times$  expected alerts at grid location (smallest square=1185, largest square=116296).  
 Colour of squares shows  $PA_{norm,g}$ , Blue = 0% PA; Red = 100% PA.



### **Population-weighted overall region measures**

Table 5-4 presents the population-weighted expected results for the entire region at the different warning radii, calculated using the Equations 4.16 to 4.18. Comparing Table 5-4 to

Table 5-2, it is clear that the population-weighted  $CTA_{norm}$  is better than the unweighted  $CTA_{norm}$ , indicating that PLUM is generating more correct timely alerts in the high population areas. The population-weighted  $PA_{norm}$  scores are higher than the unweighted  $PA_{norm}$ , but there is little difference in the  $FA_{norm}$  scores.

The population-weighted  $CTA_{norm}$  improves with increasing warning radius, from 64.4% to 73.5%, but the population-weighted  $FA_{norm}$  also increases with radius at a higher rate than the population-weighted  $CTA_{norm}$ . Therefore, it is not clear how to choose the warning radius for the best trade-off between  $CTA_{norm}$  and  $FA_{norm}$ .

*Table 5-4 Population-weighted  $CTA_{norm}$ ,  $FA_{norm}$ ,  $PA_{norm}$  for warning radii 30 km, 40 km, 50 km, 60 km and 70 km.*

<b>Warning radius</b>	<b>Population-weighted <math>CTA_{norm}</math></b>	<b>Population-weighted <math>FA_{norm}</math></b>	<b>Population-weighted <math>PA_{norm}</math></b>
30 km	64.4%	9.4%	56.4%
40 km	67.9%	13.5%	56.5%
50 km	69.6%	25.3%	51.2%
60 km	69.7%	36.2%	44.5%
70 km	72.5%	44.1%	41.5%

### **Population-weighted warning time measures**

Figure 5-18 illustrates the population-weighted mean warning times for each grid location for different warning radii. The sizes of the squares represent the total number of person-alerts for each grid location. For the significant grid locations (the 10 largest squares), the mean warning time is between -1.3 seconds and 9.8 seconds for the 30 km warning radius. For the same locations, the mean warning time increases to between -0.5 seconds and 11.7 seconds for the 50 km radius. For 70 km all significant 10 locations have positive mean warning times, ranging from 1.8 to 12.7 seconds. However, the location with the second highest expected number of correct person-alerts (the south part of Wellington City) always has the lowest mean warning time of the 10 locations.

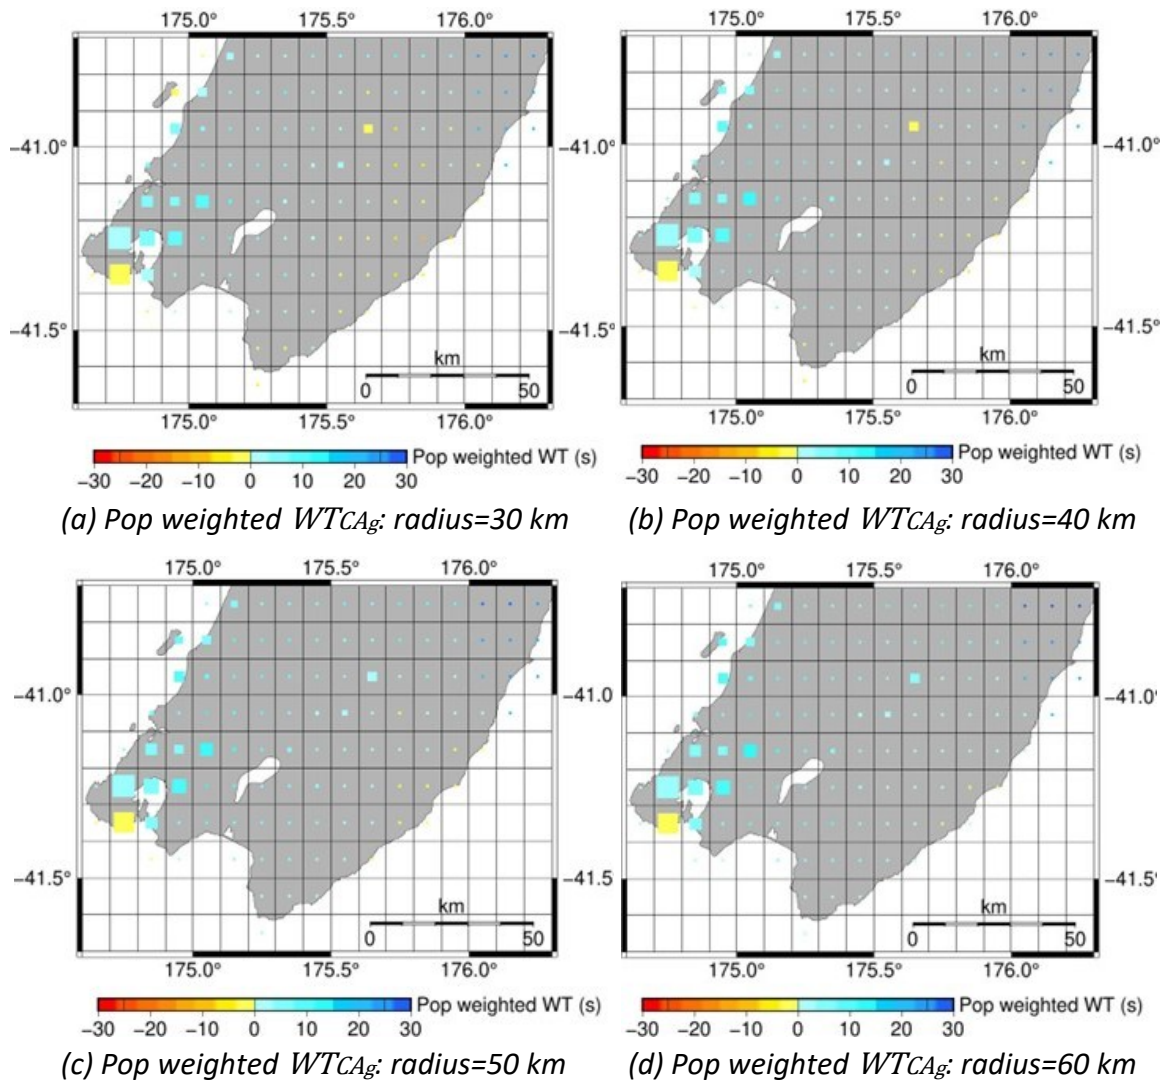


Figure 5-18 Population-weighted expected mean warning times (seconds) for all correct alerts, for warning radii (a) 30 km, (b) 40 km, (c) 50 km, (d) 60 km, (e) 70 km. Size of squares represents the expected number of correct person-alerts in the cell (smallest square=966, largest square=45687). Colour of squares represents the expected mean warning time: yellow/orange for late alerts; blue for timely alerts. Squares of size less than 1% of the largest square are shown as small circles.




Table 5-5 presents the population-weighted overall mean warning times for the entire region for all warning radii. The population-weighted expected mean warning times are always positive for all correct alerts, and which suggests that the majority of correct alerts received by the most critical locations are timely alerts. For the correct timely alerts, the warning time for the 30 km radius is 6.3 seconds which improves to 9.3 seconds for the 70 km radius.

*Table 5-5 Population-weighted mean warning times, for all correct alerts, and for correct timely alerts for warning radii 30 km, 40 km, 50 km, 60 km, and 70 km.*

<b>Warning radius</b>	<b>Population-weighted WTCTA (s)</b>	<b>Population-weighted WTCA (S)</b>
30 km	6.3 ± 2.8	3.4 ± 3.6
40 km	7.5 ± 3.2	4.4 ± 3.9
50 km	8.3 ± 3.5	5.1 ± 3.8
60 km	8.9 ± 4.0	5.6 ± 4.1
70 km	9.3 ± 3.5	6.2 ± 3.5

### **5.3.3 Analysis of alerts for south Wellington City from Cook Strait earthquakes.**

The results presented in Figure 5-15 of Section 5.3.2 indicate that high population grid cell in southern Wellington City cell (marked by the purple rectangle of Figure 5-19) experienced more missed or late alerts than correct timely alerts, even at the largest warning radius. This was the worst performing location in the region, and therefore we did further analysis to identify the cause. The hypocentres of all earthquakes that generated strong shaking for the grid cell are plotted in Figure 5-19. The red circles represent the hypocentres of all earthquakes that resulted in missing or late alerts to that grid location with a 30 km warning radius (Note, there was only one missed alert and many late alerts). The blue circles represent the earthquakes that resulted in correct timely alerts. Most of the earthquakes labelled red are located in the Cook Strait, which means there are no sensor stations that would be able to produce timely alerts to that grid cell. The few sensors on land do not detect strong shaking early enough to provide a timely alert to the grid location. Most of the earthquakes on the land or close to the South Island, where sensors are located, are able to provide correct, timely alerts to the grid location, as indicated by the blue circles on the land or close to the South Island. For the small number of earthquakes in the middle of Cook Strait for which timely alerts were generated, the fault was presumably long enough and oriented such that

shaking was able to reach one of the sensor stations in time. Even when correct timely alerts could be generated, the warning times were low: the largest warning time was 5.3 seconds, and 2/3 of the warning times were less than 2.5 seconds.

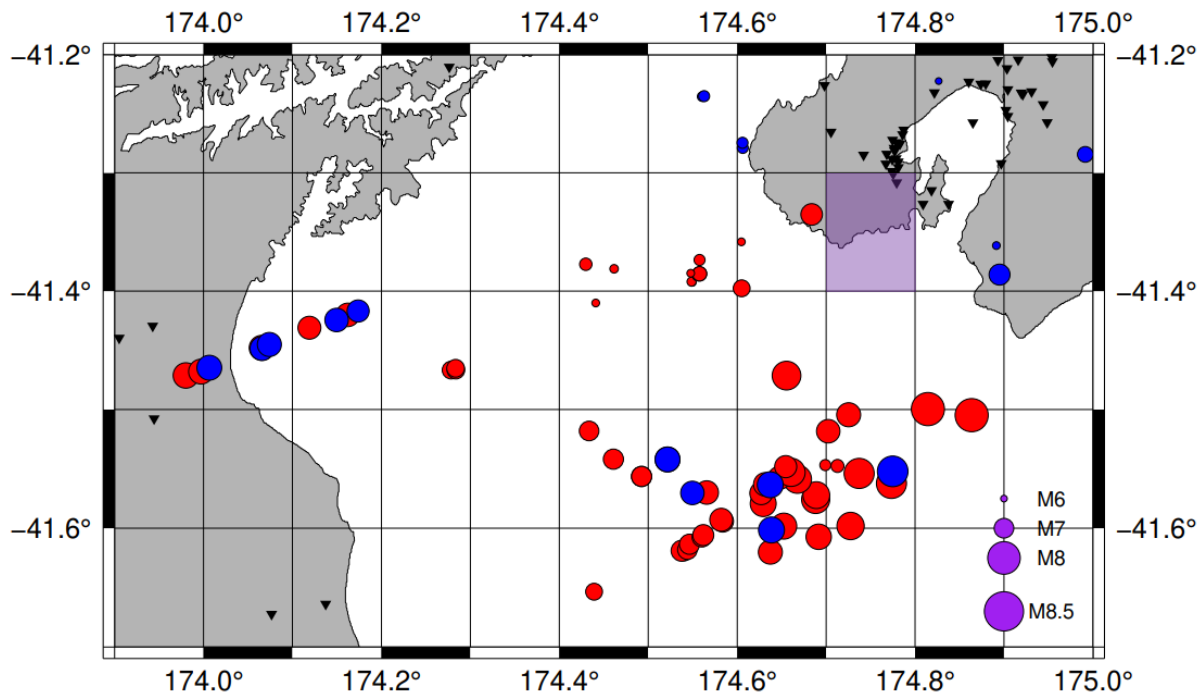


Figure 5-19 Centroids of all earthquakes causing strong shaking in south Wellington City. The purple shaded cell shows south Wellington City. Red circles are for earthquakes that resulted in late or missing alerts (at 30 km warning radius). Blue circles are for earthquakes that resulted in correct timely alerts (at 30 km warning radius). Black triangles indicate GeoNet strong motion sensor station locations.

## 5.4 Discussion

According to the results presented in Section 5.3, almost all grid locations of the Greater Wellington region are expected to experience earthquake ground motion above the MMI V during a 100-year period. However, the data shows more shaking in some of the coastal regions. One reason is that those locations are defined as site class E sites, where the ground motion amplification is high. Note that there are likely to be other small locations with more strong shaking, but this study does not consider specific amplification effects of sites such as basin effects, but only considers VS30 site classes, which have a 0.01° resolution.

The overall measures, both population-weighted (Table 5-5) and unweighted (Table 5-3), show that  $CTA_{norm}$  improves with increasing warning radius, up to 50 km, with much less

improvement at higher radii (There is an exception in the southern part of Wellington City and in Masterton, where  $CTA_{norm}$  continues to improve up to 70 km). This improvement is counterbalanced by the false alerts. The overall measures give a  $FA_{norm}$  of under 10% at 30 km warning radius, but if the warning radius is increased, even to 40 km, the false alert rate rapidly rises to levels that would be unacceptable. It is difficult to identify a warning radius that consistently leads to a good trade-off between correct timely alerts and false alerts. It may be that different parts of the region would be best served by different warning radii.

The false alert levels in the Wellington region, especially at higher warning radii, are significantly higher than the Christchurch region (as presented in Section 4.3). The underlying source of the increase in false alerts with increased warning radius is that PLUM's undamped motion assumption, which works fairly well around the 30 km radius, clearly fails when the warning radius increases, especially by 60 km or 70 km. However, this does not explain the difference between the false alert rate results for the Wellington region and the Canterbury region. A possible explanation is the greater variation in site classes in the Wellington region. A station with a D or E site class may detect strong shaking and alert a nearby grid location that has a site class of only B or C and therefore may only detect much weaker shaking. This would require a more detailed analysis that must be left to later work.

One part of the Wellington region that is expected to receive a high fraction of Correct Timely Alerts is the Kāpiti region, even at a warning radius of only 30 km or 40 km. Given the low density of sensor stations around Kāpiti, this may seem surprising. One reason for the higher  $CTA_{norm}$  is that there are very few earthquakes centred in Kāpiti itself, for which timely alerts would be impossible. The other reason is that there are a few key stations located between the Kāpiti region and the locations of the two main groupings of earthquakes to the north of Kāpiti and to the east in Wairarapa, and the location of these stations allows them to provide timely alerts to Kāpiti.

The northern part of the Wellington metropolitan area (including the north of Wellington City, Porirua and the Hutt Valley) is also expected to receive a high fraction of correct timely alerts, even at a warning radius of only 30 km. Furthermore, the fraction of false alerts is also low at 30 km. Given the high number of sensors in and around this area, the good performance is perhaps not surprising. This area covers over 60% of the population of the entire Wellington region.

For both the Kāpiti region and the northern part of the Wellington metropolitan area, the fraction of false alerts rapidly rises to problematic levels for higher warning radii and for both areas, the most appropriate choice of warning radius might be 30 km.

Generally, the southern part of Wellington City (10% of the total population of the region) and the central Wairarapa region including the regional population centre of Masterton (4% of the total population of the region) had poor performance from the PLUM algorithm. In particular, they experienced more missed or late alerts than correct timely alerts for all warning radii and the warning times were low. The maximum warning time for a timely alert in the southern part of Wellington City was less than the median warning time for timely alerts across the whole region. An important reason is the relative locations of the earthquake events and the sensor stations around those regions. According to Figure 5-6, the catalogue contains many large earthquakes in Cook Strait, to the south of Wellington City. Section 5.3.3 showed that those earthquakes were the main source of strong shaking events at the grid location in the southern part of Wellington City, and that most of those earthquakes resulted in late alerts rather than timely alerts. As shown in Figure 5-19, there are many stations in the Wellington metropolitan region, but very few stations along the south coast near Wellington City or at the south of Wairarapa. Therefore, there are no stations between the earthquakes in Cook Strait and the southern part of Wellington City that could provide timely warnings. Similarly, there are also no stations surrounding Masterton that could provide timely warnings for the earthquakes in central Wairarapa.

The PLUM algorithm works well for Christchurch but not as well for Wellington City, especially for earthquakes in Cook Strait and the Hikurangi subduction zone. Increasing the warning radius might enable us to obtain alerts from the South Island and further Wairarapa sensors. However, just increasing the warning radius is not the best solution as it also increases the occurrence of false and precautionary alerts, and the more distant stations are still unable to provide timely alerts with sufficient warning time.

The key problem is that there are no GeoNet stations in Cook Strait and an insufficient number in southern Wairarapa, which means the EEW system cannot effectively provide warnings for earthquakes in these areas, particularly along the Hikurangi subduction zone and Cook Strait. This lack of coverage is a significant issue because these regions have many earthquakes, and

without proper coverage, Wellington City, especially the southern part, will miss out on correct timely alerts for earthquakes from these sources.

A better approach than increasing the warning radius would be placing a few more sensors in key areas instead of just expanding the warning radius. Additional stations placed along the southern coast would help provide timely alerts to Wellington City, but warning times may remain short since the stations would be so close to the population centre. Adding offshore sensors in Cook Strait would also be beneficial, but this is more challenging due to the strong currents, variable ocean conditions, and the depth and ruggedness of the seafloor (Stevens & Smith, 2009). Additionally, harsh weather, high winds, and waves further complicate sensor deployment and maintenance. Data transmission of real-time offshore sensors is also more complex as it involves underwater cables or wireless communication systems and communication under the water within the dynamic environment of Cook Strait. However, Japan has offshore seismic monitoring networks (Dhakal & Kunugi, 2021). Therefore, the feasibility and cost-benefit analysis to install offshore sensors should be conducted in a future study considering the EEW performance improvement for the possible installation locations.

For the Masterton area, adding a few stations around Masterton (about 30 km from the city) would also greatly improve the warning performance for that area.

With additional sensors in key locations, the PLUM algorithm would be able to provide reasonable EEW for the Wellington region.



# Chapter 6 Conclusions and Future Work

This chapter presents the key conclusions of our testing PLUM in the Canterbury and Wellington regions and provides recommendations based on the results. It also identifies some of the limitations of the current study and suggests areas of potential future work that build on the thesis.

## 6.1 Discussion and recommendations

### Latency of the GeoNet sensor network

The communication latency analysis of the GeoNet strong motion sensor network (Section 3.1) showed that the current latency in the sensor network is around 4-5 seconds. However, the current network was designed for earthquake location detection in order to build catalogues rather than EEW. The analysis in this thesis considers that a 2 second latency, including communication delay, is feasible by changing the packetization technique of the data acquisition systems and enhancing the processing in the station to include pre-processing such as PLUM threshold detection. The changes would reduce the number of data samples in each packet (sending more but smaller packets) and would require the stations to immediately send very small packets, just for EEW purposes, whenever the PGA went above critical thresholds. These modifications could be implemented without replacing the hardware to cater to the EEW requirements. We recommend that this kind of EEW specific modifications to the GeoNet sensor network should be considered in future EEW system design process, particularly if the PLUM algorithm is to be used. However, modifying data packetisation could reduce compression efficiency, increasing telemetry bandwidth requirements and costs. Therefore, this should be done with careful consideration of trade-offs to achieve the best efficiency within infrastructure limitations.

Stubailo et al. (2020) highlight that high bandwidth is required during elevated shaking of large magnitude events, making packet compression less effective. To address this, it is important to consider either compressing techniques at the datalogger or reducing the amount of data sent to the server through local preprocessing.

### **Using a synthetic earthquake catalogue to evaluate an EEW system.**

The thesis provides an example of how to use a physics-based synthetic earthquake catalogue to evaluate the expected performance of an EEW algorithm on future possible seismic events. Evaluating an algorithm on a single earthquake event (e.g., an Alpine fault earthquake scenario) is problematic because the algorithm may work well for the one event but may not work well for other earthquakes and the results cannot be generalised to future earthquakes. Evaluating an algorithm on recorded data from actual earthquakes is better, but is limited because of the small amount of recorded data, especially of larger earthquakes. The thesis showed how to use the RNC2 RSQSim catalogue (Shaw et al., 2022) to determine the expected performance of an EEW algorithm on a range of different earthquake events that could affect the Canterbury and Wellington regions over the coming 100-year period (Chapters 4 and 5). This approach gives much more generalisable results than only using recorded data, and would be appropriate to use for evaluating other EEW algorithms.

There were some limitations in the use of the catalogue for this analysis. Most importantly, the catalogue did not include any events due to blind unmapped faults, for example, the blind fault under Christchurch City, which caused the very damaging Christchurch earthquake. We used the first stable and published version of the synthetic earthquake catalogue, as the more recent developments in the catalogue were being conducted in parallel to the research in this thesis. Future research should use the latest, updated version of the catalogue, which will address some of the limitations of the first catalogue. If the catalogue does not include events due to blind/unknown/unmapped faults, it would be important to add some appropriately generated random events that could at least somewhat account for events due to such faults (e.g., through a distributed seismicity model; Gerstenberger et al, 2024).

### **Simulating ground motion of synthetic earthquakes with EXSIM**

The EXSIM validation analysis, which was performed using both generated data and recorded data of Darfield earthquake (Section 3.3), showed that the EXSIM ground motion simulator is an appropriate tool to generate ground motions for this kind of research. This is important for two reasons. First, we need to test EEW algorithms on simulated earthquakes in order to provide a large enough set of test data. Second, evaluating an EEW algorithm requires earthquake ground shaking data for many locations where there is no recorded data from

actual earthquakes, so that even if the evaluation uses historical data, there will still be a need to use a ground motion simulator to evaluate an EEW algorithm effectively.

There are some significant limitations to EXSIM. The first is that it only generates ground-motion of S-waves, and does not simulate P-waves. The limitation is not a major problem for the PLUM algorithm since PLUM is designed to use S-waves. However, some other EEW algorithms, such as EPIC (Chung et al., 2019) also depend on P-waves. Evaluating such algorithms would require a different simulator such as SPECFEM3D (Komatitsch & Tromp, 1999; Peter et al., 2011). SPECFEM3D uses a spectral element method and is much more computationally expensive than EXSIM. Furthermore, PLUM depends on detecting strong shaking, and the P-waves of large earthquakes could generate strong shaking that PLUM could use to provide earlier warnings. Using a simulator other than EXSIM would be required to evaluate this. Moreover, EXSIM utilizes a 1D attenuation and geological model, which limits its ability to handle complex, geologically diverse regions like the Taupō volcanic zone.

### **PLUM Results**

PLUM performed well for the Canterbury region, particularly for the high population regions around Christchurch City (Chapter 4). In the high population areas around Christchurch City, it provided timely alerts for at least 94% of the strong shaking events and generated a very small fraction of false alarms (see Figure 4-17 and Table 4-6). Mean warning time of correct alerts for the centre of Christchurch City is above 11 seconds for all warning radii. The performance was not as good in the western part of the region, with a higher fraction of missed or late alerts, and lower warning times, particularly at the minimum warning radius of 30 km. Across the region, increasing the warning radius improved the fraction of timely alerts, but also increased the false alert numbers. A 50 km warning radius appeared to give a good trade-off of timely alerts vs false alerts.

PLUM did not work as well for the Wellington region, although it still performed reasonably well for certain areas, especially the northern part of the Wellington metropolitan area, Kāpiti, and the northeast corner of the region (Chapter 5). It performed particularly poorly for the southern part of Wellington City and for most of the Wairarapa, including Masterton. Overall, the fraction of timely alerts for strong shaking events was a bit lower than in Canterbury, with lower warning times, and the false alarm rate was very much higher than in Canterbury. As in

Canterbury, increasing the warning radius improves the fraction of timely alerts, but even at 70 km, much of the Wairarapa had correct timely alerts for less than 50% of the strong shaking events. Furthermore, the fraction of false alerts increased rapidly as the warning radius expanded. It is more difficult to identify a warning radius giving a good trade-off between timely alerts and false alerts.

It is clear that the sensor locations and density are crucial to the performance of the PLUM algorithm. To provide a warning to a specific location for an earthquake, PLUM requires a sensor that is closer to the earthquake than to the location and close enough to the location to issue a warning. The areas with good performance, particularly Christchurch, had a reasonable number of sensors around the areas, positioned between the areas and the locations of most of the earthquakes. The areas with poor performance did not have sufficient sensors around them. Notably, the southern part of Wellington City, which has a high population (10% of the population of the entire Wellington region), experienced particularly poor performance with a very low fraction of correct timely alerts (always less than 50%  $CTA_{norm}$ , see Figure 5-15). Almost all of the earthquakes that caused strong shaking in Wellington City were located in Cook Strait (see Figure 5-19), and there are no sensors located between them and the southern part of Wellington City, so it was not possible for PLUM to provide timely warnings (see Section 5.3.3). Other areas with poor performance, such as the western half of the Canterbury and the Wairarapa, also had a very low density of stations.

The most important recommendation for implementing a useful EEW system is to increase the number of sensors in these critical areas. The analysis in Section 3.4.2 presented an effective method for determining good locations for additional sensors, which used EEW performance to govern the search for the best sensor locations. This method would be useful for western Canterbury and the Wairarapa. Adding sensors to provide effective warnings to Wellington City would be more challenging. The first step would be to install strong motion sensors in existing GeoNet stations in South Wairarapa and south of Wainuiomata. The second step would be to install new sensor stations along the South Coast of Wellington and Wairarapa. These could provide timely alerts for many of the Cook Strait earthquakes, but the warning times would be short. Note that timely alerts with very short warning times may not be very useful for warning the general public but may still be useful for site-specific EEW that allow automated actions for industrial or other large sites, for example, to shut down

equipment, or turn off water, electricity and gas. A more effective, though more expensive, step would be to install sensors such as ocean bottom seismometers in Cook Strait itself. There are a variety of recent technologies for this which may become feasible, for example using undersea SMART fiber optic cables (Wallace, 2021) or Distributed Acoustic Sensing (DAS) using existing fiber optic cables (Farghal et al., 2022). Another potential low-cost solution to consider is deploying buoys/ floats (Sukhovich et al., 2015) in Cook Strait to detect P-waves. This approach could serve as an alternative to installing ocean bottom seismometers and may also be effective in triggering the PLUM algorithm for large earthquakes.

In both Canterbury and Wellington, the fraction of precautionary alerts (where an alert is received, but only weak shaking occurs) is high. A high frequency of precautionary alerts has the potential of leading to “alert fatigue” where people begin ignoring warnings if they receive too many alerts that are not followed by a dangerous event (Becker et al., 2020a; Vinnell et al., 2023). However, the frequency of alerts (on average, less than 1 every 5 years) is not high enough to lead to alert fatigue, though it could be a problem during an aftershock sequence of a large earthquake. Changing the alert threshold to a higher value during an aftershock sequence may be a solution to avoid alert fatigue in this particular situation. Alternatively, the system could allow users to set parameters specifying which alerts they wanted to receive.

Section 4.2.7 and 4.2.8 introduce performance metrics we used to assess PLUM in Canterbury and Wellington. These metrics provide insights into the effectiveness of PLUM which can inform stakeholders such as the National Emergency Management Agency (NEMA) and the Natural Hazards Commission (NHC). They would need to consider factors such as the acceptable percentages of false and precautionary alerts, as well as the required accuracy and threshold levels for the alerts.

While the current study focuses on magnitude 6 and above earthquakes, an operational PLUM EEW system would face many smaller events that could possibly exceed trigger thresholds. These smaller events can also contribute to higher precautionary and false alert rates. Real-world deployments will need to raise detection and confirmation thresholds to filter out low-impact events. Tuning these parameters is needed for balancing timely warnings against alert fatigue and for maintaining acceptable false-alert fractions.

A limitation of the analysis is that each grid cell was assigned a single site class, even when there were a range of site classes across the grid cell. A consequence is that the analysis may have underestimated the shaking that would be experienced in some parts of a grid cell. For example, the 2016 Kaikōura Earthquake caused building damage in some parts of Wellington City that are site class E (Orense et al., 2017) though the shaking in other parts of the city was not as strong. The analysis assumed that grid cells covering Wellington City were site class C, and therefore would not have considered that those cells were needed for an alert for the Kaikōura Earthquake. Possible future work that could address this limitation is discussed in Section 6.2.

### **Limitations of PLUM**

While the PLUM algorithm worked well in Christchurch and showed potential in parts of Wellington, there are several inherent limitations to be considered.

One major limitation is that PLUM can only give warnings up to a certain maximum time, which is controlled by the warning radius setting. This radius sets how far from the shaking the system will send alerts. Even if shaking is detected early, PLUM won't send an alert to areas outside that set radius. As a result, the maximum possible warning time is limited by the chosen radius, regardless of earthquake magnitude or event characteristics.

In comparison, some early warning algorithms that try to estimate the earthquake source can sometimes give much longer warnings, especially for large and distant earthquakes. This limitation is linked to the trade-off discussed earlier in Sections 4.4 and 5.4—increasing the warning radius improves the likelihood of issuing timely alerts but simultaneously increases the risk of precautionary or false alerts, especially in areas with low sensor density.

Another limitation is how PLUM detects an earthquake. It needs a certain number of stations to detect and confirm it—usually two stations are enough to trigger an alert. This makes PLUM fast and useful in places where there aren't many sensors. However, PLUM uses higher ground motion thresholds than some other algorithms, especially those based on P-waves. These other methods can sometimes detect earthquakes a bit earlier. Because PLUM waits for stronger shaking, it might miss the very first signs of an event, which could lead to a slightly delayed alert, especially for locations distant from the source.

PLUM is based on the key assumption that PGA does not attenuate as seismic waveforms propagate. This allows the algorithm to remain simple by not requiring the identification of the earthquake source. But shaking usually weakens over distance, and this assumption starts to break down, especially when larger warning radii are used. That means PLUM's accuracy drops as the warning radius increases. PLUM also does not handle variations like source directivity, which can cause shaking to be stronger in certain directions, or regional geological variations.

Overall, while PLUM offers simplicity, speed, and reasonable robustness—especially in areas with good sensor coverage—it is important to consider the algorithm's limitations. These limitations need to be considered when planning or reviewing how PLUM is used in operational EEW systems. Future work should look into combining PLUM with other EEW approaches, improving how local site effects are handled, and using better ground motion models to address some of these challenges.

### **PLUM site corrections**

In our analysis, site corrections were added only to the synthetic seismogram data, not to the PLUM predictions, and we follow a similar approach to that used on the US West Coast by Cochran et al. (2019). We cannot do site corrections like Japan does (Iwakiri et al., 2011; Kodera et al., 2018) because we do not have the site correction data for the PLUM algorithm that they have to estimate site effects at each target location based on observation sites.

Applying site corrections in public EEW systems is also challenging because of high uncertainties. It is difficult to define accurate corrections for all target sites. So, in this study, EEW alerts are issued without considering target site corrections, assuming uniform site conditions across all locations (however, we consider the site class of each station and grid location while evaluating the performance, as we incorporate site classes while generating synthetic seismograms). Site-specific corrections could be useful if we are alerting a single place, like a building or a factory. But for public EEW, using a common ground model is more practical.

Cochran et al. (2019) tested PLUM on the US West Coast and did not apply site corrections to the observed or predicted intensity values. They reported that site corrections had only a

small impact on the results—less than 15% change. They also mentioned that VS30-based site corrections only do a moderate job because they represent the average response and cannot capture outlier amplification at some sites. However, future work could look into better ways to include site corrections to improve EEW accuracy.

### **Failure scenarios**

Sections 4.2.9 and 4.3.3 explored the consequences of sensor failure on the performance of the PLUM algorithm for Canterbury. The results showed that neither of the two different failure scenarios had a major impact on the performance of the PLUM algorithm. Losing 33% of all sensors caused only a 3% reduction in the fraction of correct timely alerts in the higher population areas of Christchurch City and only a small reduction in warning times. There was a little more impact from losing 33% of the sensors in western Canterbury where there is a low sensor density. Losing all the sensors in a band around Christchurch had a larger effect on Christchurch (up to 11% reduction) but the network was shown to be surprisingly robust against quite substantial sensor failure. Adding stations in areas of low sensor density would help increase the resilience of the network, as would ensuring that the components of the sensor and communication network are sufficiently independent that large scale concurrent failure of a large number of the sensors in an area is very unlikely.

The results with a 33% station failure (representing disconnection of randomly chosen sensors) suggest that using different service providers is important for the robustness, but it is important that the stations should be distributed among the service providers in a way that avoids groups of all stations in one area being assigned to a single provider. Greater robustness could be achieved by maintaining communication redundancy through the use of more than one communication link/service provider for the same station location.

## **6.2 Future work**

In addition to the suggestions above for mitigating the limitations, there are a number of directions for future work which could improve or build on the results of this thesis.

## Defining the grid cells

In this study, the regions were divided into  $0.1^\circ \times 0.1^\circ$  grid cells and we assumed that each grid cell would receive an alert separately. A problem with this approach is that the grid cells are not aligned with geographical features or with the location of the population, so that the simulation has to take averages of the site classes across the grid cell, and the alerts do not take into account where the people receiving the alerts are located.

This grid size is finer than the  $0.2^\circ \times 0.2^\circ$  grid cells used in the US shake alert EEW system and the PLUM analysis on the US West Coast (Cochran et al, 2022) and therefore has to make less approximations than those systems. Japan's EEW is organized into 188 regions called seismic hazard blocks, with each prefecture containing approximately four on average (Fujinawa & Noda, 2013), which may enable more homogeneous cells. Minson et al. (2020) also analyse PLUM in California using an approach similar to Japan's regional block approach. They considered boundaries of California counties, the National Weather Service forecast zones as well as using a 50 km grid based approach.

However, a better approach than using arbitrary geographical boundaries would be to partition the region according to the alert dissemination technology, for example, using cell tower locations to define the grid cells (Cardone et al., 2014). This would mean that the performance scores would better reflect the more realistic real world conditions. If the alert dissemination technology was also able to use GPS data from mobile devices to target alerts even more finely (McBride et al, 2023), then utilising an even finer grid size might be achievable.

The VS30 site classification model used a  $0.01^\circ \times 0.01^\circ$  grid. In the Wellington region in particular, there was considerable variation in the VS30 site classes across many of the  $0.1^\circ \times 0.1^\circ$  grid cells. It would be possible to use the finer  $0.01^\circ \times 0.01^\circ$  grid for the PLUM analysis and obtain much more accurate measures of EEW performance. Note that this would require simulating the ground motion at 100 times more grid locations, which would increase the analysis time considerably.

## **Modifying the PLUM algorithm**

PLUM was originally designed to have a warning radius of 30 km since this is compatible with the undamped propagation assumption that underlies the design of PLUM. The density of sensor stations in Japan is high enough that this is a reasonable warning radius. For New Zealand, where the sensor density is much lower in some areas, it is necessary to increase the warning radius to enable PLUM to be able to issue warnings to most locations. The higher warning radius is also necessary to allow longer warning times. However, a larger warning radius starts to violate the undamped propagation assumption, and this is one reason for the high number of precautionary alerts and false alerts. The results of the analysis in the thesis suggested that a warning radius of at least 40 km to 50 km might be an appropriate compromise in the NZ context.

To enable enough alerts to be issued and to improve the warning time for distant large earthquakes without increasing the precautionary and false alert rate, PLUM needs to be modified to take into account the attenuation in the propagation. A simple way of doing this is to require a higher shaking threshold at the alerting station when issuing a warning to a more distant location. An obvious future work would be to repeat the analysis of this thesis using multiple shaking thresholds and issuing alerts to more distant locations for the higher shaking events. Saunders et al. (2024) have very recently proposed an algorithm called APPLES (Attenuated ProPagation of Local Earthquake Shaking) for the US West Coast version of the PLUM that similarly relaxes the undamped propagation assumption of PLUM.

## **Confirmation and noise mitigation strategies.**

The alert confirmation strategy used in the thesis requires verification from another station to ensure that detected shaking is valid rather than noise. Where the station density is low (as it is in many rural areas in NZ), there are few stations that can provide the verification, and their distance can delay an alert, and in some cases, the lack of stations may even prevent an alert from being issued. It is possible that this limitation could be reduced by incorporating one of the other verification strategies indicated in Kilb et al. (2021), for example, confirming an alert by a co-located velocity sensor or by another complementary type of sensor located in the same place. These strategies would remove the need for confirmation from another station in regions of low sensor density. Co-located sensors can help reduce instrumental

noise by allowing the system to confirm a signal from different instruments at the same site. However, they may increase processing time and add complexity to the system. Future work could explore whether this would significantly improve the performance of PLUM in the NZ context and possibly provide recommendations to GNS to co-locate complementary sensors in some existing sensor stations.

The analysis in the thesis did not use any noise in the ground motion data. Future work would need to explore the performance of PLUM with the alert confirmation strategies when the ground motion data contains real noise. This might involve testing PLUM with real, continuous data, or adding realistic simulated noise to simulated ground motion data. The effect of multiple earthquakes, such as swarms or aftershock sequences, also needs to be considered. As an example, subduction zone earthquake could activate nearby crustal faults as well (Hicks & Rietbrock, 2015). The robustness of the PLUM algorithm should be tested under such complex scenarios.

### **Using recorded data from small earthquakes**

The thesis used a catalogue of synthetic earthquakes to evaluate the PLUM algorithm because there is not sufficient historical data on earthquakes at magnitude 6 and above to provide a reasonable evaluation of an EEW algorithm. However, Petersen et al. (2015), mention that small earthquakes usually occur in areas where larger earthquakes are also likely to happen. Based on that assumption, it may be feasible to test PLUM on historical data of smaller earthquakes for which more recorded data is available. By setting appropriate lower thresholds for both detection and evaluation, the performance could be assessed and then could be generalised to understand how the algorithm might perform on larger earthquakes. Future work could look at using historical data of smaller earthquakes in this way to complement the analysis using synthetic earthquakes.

Artificial neural networks (ANN) could also be explored as part of future work. So, we could train an ANN on station data, real or synthetic, and use it to detect shaking patterns and issue alerts—like a smarter version of PLUM. With AI and machine learning constantly improving, there is a lot of potential to make EEW faster and more accurate. One of the challenges will be obtaining enough high-quality data from real earthquake scenarios to validate the system, which is currently limited.

### **Quantify directivity impacts**

In the analysis of Chapter 4 and 5, the vertical-fault assumption affects shaking intensities very close to the fault by reducing hanging-wall amplification and increasing footwall shaking (Donahue & Abrahamson, 2014). Assuming bilateral rupture produces median estimates of ground motion without the extremes introduced by unilateral directivity pulses. However, for large events, fault geometry and rupture directivity can amplify motions in one direction; by using bilateral faults, these directivity effects are not simulated, so the simulations may underestimate peak motions on one side and overestimate them on the other. These effects will cancel out on average when considering statistical effects over many 100-year periods.

Therefore, it will be interesting to explore the effect of directivity, as well as more variable fault dips in future work—this is especially relevant for faults likely prone to unilateral propagation (e.g., Alpine Fault unilateral vs. bilateral scenarios).

### **Evaluate PLUM with earthquakes from small to large magnitudes and under varying threshold settings.**

The main Canterbury and Wellington studies only consider earthquakes of M 6.0 and above. However, a small event could also generate some alerts, including false and precautionary ones. Thus, testing PLUM across the full magnitude range and under different detection and confirmation thresholds would be useful. By adjusting these thresholds, we can filter out low-impact events and fine-tune the balance between unnecessary alerts and timely warnings. Evaluating PLUM with earthquakes from small to large magnitudes and under varying threshold settings to properly assess these trade-offs would be an interesting future task.

### **Extending PLUM to generate target-based alerts**

The PLUM algorithm used in this thesis detects strong ground motion at a sensor and issues early warnings to all target locations within the warning radius (e.g., within 30 km of the sensor). This approach ignores site-specific factors, assuming that all target locations within the warning radius will experience the same shaking intensity as the sensor. This means that some stations get false alerts or precautionary alerts, and other have missing alerts.

It would be better if the algorithm took into account site-specific factors of the target location in order to provide target-based alerts. For example, the relative site classes of the sensor and the target should be taken into account, so a sensor with site class C should detect shaking at a lower threshold when sending an alert to a target with site class E, but detect shaking at a higher threshold when sending an alert to a target with site class A. Site-specific information such as basin effects could also be used to target alerts more appropriately. Some locations with specific needs, such as industrial sites or power plants, might also need tailored alerts.

Future work should extend the PLUM algorithm to detect shaking at sensor stations at a range of different thresholds, and then issue alerts to different target locations at different levels, taking into account the site-specific factors of the sensor and targets. Note that this could also be used to take attenuation into account, issuing alerts to more distant sites only at higher shaking thresholds and enabling longer warning times.

#### **Combining with complementary algorithms and EEW systems.**

Future work should consider how a GeoNet based PLUM EEW could be combined and integrated with other EEW algorithms and other systems. Two projects that ran concurrent with this thesis should certainly be considered in this way. Prasanna et al. (2022) proposes an EEW system using low-cost sensors such as Raspberry Shakes. That kind of low-cost sensor network could be useful to complement the GeoNet high-end network by filling the gaps where GeoNet sensor density is low. Integrating with low-cost sensor networks can improve station density, which helps with faster and more accurate event detection. But they may introduce more noise into the system due to lower sensor quality. However, with proper noise mitigation techniques or AI-based filtering, it may still be useful to include low-cost sensors in the network.

Combining with the data of Google Android earthquake detection, which uses mobile phone accelerometers to detect shaking, would be another area to consider. However, limited information is available about that algorithm as it is a proprietary system.

The FinDer algorithm (see Section 2.1.2), implemented by Andrews et al., (2024) is a useful tool for rapid responders and scientists to identify the characteristics of an earthquake. However, it can take minutes to detect and deliver the information about an earthquake. It

would be a useful complementary algorithm to detect and issue EEW for high magnitude and complex earthquake scenarios, as it could provide warning to larger distances than the PLUM limited warning radii.

It is also very important that any EEW system for NZ, whether using PLUM or other algorithms, is integrated with the NZ tsunami warning system (Gusman et al., 2022; Lacoua et al., 2024). Future work could explore how that integration could be done. Additionally, the EEW system should not interfere with the "long or strong, get gone" tsunami self-evacuation procedure (Vinnell et al, 2022) of New Zealand.

### **Network optimisation**

A critical aspect of any network-based EEW system is the rapid, reliable transmission of sensor data. Achieving this requires refining existing transmission practices and implementing methods to handle network disruptions during extreme events. One suggestion as a future work to reduce latency is to modify the packetization techniques within the data acquisition systems. Decreasing the number of samples per packet and introducing local processing to reduce network traffic could transmit the data much more quickly. In addition to that, future work could also address network resilience and failure mitigation.

Some of the future work suggestions are:

- Simulate network failures: Develop simulations to test how the network responds to packet loss, sudden latency spikes, sudden traffic during large events, and link failures, and design methods to transmit data efficiently during critical moments.
- Network redundancy: Explore how to use redundant communication channels for EEW, if the primary network becomes congested or fails.
- Detailed analysis of each latency component within the network. The total latency comprises several parts including:
  - Packetisation time: Currently the GeoNet network is designed for data recordings, this can be changed to support immediate transmission for EEW.
  - Transmission time on the communications network: This is a significant component of the data latency. Conducting real-time experiments will help in finding out the variations across different sensors and regions.

- Other network components: An analysis that breaks down other network components of EEW (server processing time, alert delivery time *etc.*) will identify where further optimisations are possible.
- Sensor and data logger variability: Evaluating how different sensor types and data loggers influence overall latency, including studies between rural and urban stations.
- Communication link types: Assessing performance differences of various communication channels (cellular, satellite, dedicated radio links, *etc.*)
- Packetisation strategies: Experimenting with dynamic packet sizing or dual-mode transmission strategies—one for regular monitoring and another for EEW.

The continual advancements in communication technology also offer exciting possibilities. How to use emerging high-bandwidth, low-latency communication protocols for EEW will be an interesting future work.

## References

- Abrahamson, N. A., Silva, W. J., & Kamai, R. (2014). Summary of the ASK14 ground motion relation for active crustal regions. *Earthquake Spectra*, *30*(3), 1025-1055.
- Alessi, F., Massini, A., & Basili, R. (2012). Accelerating the production of synthetic seismograms by a multicore processor cluster with multiple GPUs. In *2012 20th EUROMICRO International Conference on Parallel, Distributed and Network-Based Processing* (pp. 434–441).
- Allen, R. M., & Kanamori, H. (2003). The potential for earthquake early warning in southern California. *Science*, *300*(5620), 786–789.
- Allen, R. M., & Melgar, D. (2019). Earthquake early warning: Advances, scientific challenges, and societal needs. *Annual Review of Earth and Planetary Sciences*, *47*, 361–388.
- Allen, R. M., Brown, H., Hellweg, M., Khainovski, O., Lombard, P., & Neuhauser, D. (2009). Real-time earthquake detection and hazard assessment by Elarms across California. *Geophysical Research Letters*, *36*(5).
- Ancheta, T. D., Darragh, R. B., Stewart, J. P., & Seyhan, E. (2013). *PEER NGA-West2 Database*. Pacific Earthquake Engineering Research Center.  
[https://peer.berkeley.edu/sites/default/files/webpeer-2013-03-timothy\\_d.\\_ancheta\\_robert\\_b.\\_darragh\\_jonathan\\_p.\\_stewart\\_emel\\_seyhan.pdf](https://peer.berkeley.edu/sites/default/files/webpeer-2013-03-timothy_d._ancheta_robert_b._darragh_jonathan_p._stewart_emel_seyhan.pdf)
- Andrews, J., Behr, Y., Böse, M., Massin, F., Kaiser, A., & Fry, B. (2024). Rapid Earthquake Rupture Characterization for New Zealand Using the FinDer Algorithm. *Bulletin of the Seismological Society of America*, *114*(2), 775-793.

- Atkinson, G. M., & Assatourians, K. (2015). Implementation and validation of EXSIM (a stochastic finite-fault ground-motion simulation algorithm) on the SCEC broadband platform. *Seismological Research Letters*, *86*(1), 48–60.
- AWS (2021), *AWS case study: GeoNet*, Retrieved from <https://aws.amazon.com/solutions/case-studies/geonet>
- Ballance, P. F. (2017). *New Zealand geology: An illustrated guide*. Geoscience Society of New Zealand Miscellaneous Publication.
- Beavan, J., Motagh, M., Fielding, E. J., Donnelly, N., & Collett, D. (2012). Fault slip models of the 2010–2011 Canterbury, New Zealand, earthquakes from geodetic data and observations of postseismic ground deformation. *New Zealand Journal of Geology and Geophysics*, *55*(3), 207–221.
- Becker, J. S., Potter, S. H., Prasanna, R., Tan, M. L., Payne, B. A., Holden, C., ... Johnston, D. M. (2020a). Scoping the potential for earthquake early warning in Aotearoa New Zealand: a sectoral analysis of perceived benefits and challenges. *International Journal of Disaster Risk Reduction*, *51*, 101765.
- Becker, J. S., Potter, S. H., Vinnell, L. J., Nakayachi, K., McBride, S. K., & Johnston, D. M. (2020b). Earthquake early warning in Aotearoa New Zealand: a survey of public perspectives to guide warning system development. *Humanities and Social Sciences Communications*, *7*(1), 1–12.
- Berryman, K. (2012). Geoscience as a component of response and recovery from the Canterbury earthquake sequence of 2010–2011. *New Zealand Journal of Geology and Geophysics*, *55*(3), 313–319. <https://doi.org/10.1080/00288306.2012.702674>
- Blake, D., Johnston, D., Leonard, G., McLaren, L., & Becker, J. (2018). A citizen science initiative to understand community response to the Kaikoura earthquake and tsunami warning in Petone and Eastbourne, Wellington, Aotearoa New Zealand. *Bulletin of the Seismological Society of America*, *108*(3B), 1807–1817. Scoping
- Boore, D. M. (2005). SMSIM—Fortran programs for simulating ground motions from earthquakes: version 2.3—a revision of 96–80-a. *US Geological Survey open-file report, US Geological Survey open-file report 00–509, revised*, *15*, 55.
- Böse, M., Andrews, J., Hartog, R., & Felizardo, C. (2023). Performance and next-generation development of the finite-fault rupture detector (FinDer) within the United States West Coast ShakeAlert warning system. *Bulletin of the Seismological Society of America*, *113*(2), 648–663.
- Böse, M., Andrews, J., O'Rourke, C., Kilb, D., Lux, A., Bunn, J., & McGuire, J. (2022a). Testing the ShakeAlert Earthquake early warning System using synthesized earthquake sequences. *Seismological Research Letters*, *94*(1), 243–259. <https://doi.org/10.1785/0220220088>
- Böse, M., Heaton, T. H., & Hauksson, E. (2012). Real-time finite fault rupture detector (FinDer) for large earthquakes. *Geophysical Journal International*, *191*(2), 803–812.
- Böse, M., Papadopoulos, A. N., Danciu, L., Clinton, J. F., & Wiemer, S. (2022b). Loss-based performance assessment and seismic network optimization for earthquake early warning. *Bulletin of the Seismological Society of America*, *112*(3), 1662–1677.

- Böse, M., Smith, D. E., Felizardo, C., Meier, M. A., Heaton, T. H., & Clinton, J. F. (2018). Finder v. 2: Improved real-time ground-motion predictions for m2–m9 with seismic finite-source characterization. *Geophysical Journal International*, 212(1), 725–742.
- Brake, K. (2018). *Youth perspectives on participation and inclusion in city life post-disaster: A Christchurch case study* (Master's thesis, University of Canterbury). Retrieved from <https://ir.canterbury.ac.nz/server/api/core/bitstreams/16d3f561-07e5-4ce4-aca8-c0007fa7f0b1/content>
- Cardone, G., Cirri, A., Corradi, A., Foschini, L., Ianniello, R., & Montanari, R. (2014). Crowdsensing in urban areas for city-scale mass gathering management: Geofencing and activity recognition. *IEEE Sensors Journal*, 14(12), 4185-4195.
- Carey, J. M., Mountjoy, J. J., Crutchley, G. J., Petley, D. N., Holden, C. F., Kaneko, Y., & Huhn, K. (2022). Episodic movement of a submarine landslide complex driven by dynamic loading during earthquakes. *Geomorphology*, 408, 108247.
- Caruso, A., Colombelli, S., Elia, L., Picozzi, M., & Zollo, A. (2017). An on-site alert level early warning system for Italy. *Journal of Geophysical Research: Solid Earth*, 122(3), 2106-2118.
- CDEM. (2002). *Civil Defence Emergency Management Act 2002, Public Act 2002 No 33*. New Zealand Government. Retrieved from <http://www.legislation.govt.nz/act/public/2002/0033/latest/DLM149789.html>
- Chandrakumar, C., Tan, M. L., Holden, C., Stephens, M. T., & Prasanna, R. (2024, May). Adapting PLUM: Earthquake early warning with node-level processing in New Zealand. In *Proceedings of the International ISCRAM Conference*.
- Chandrakumar, C., Tan, M. L., Holden, C., Stephens, M. T., & Prasanna, R. (2025). Performance analysis of P-wave detection algorithms for a community-engaged earthquake early warning system—a case study of the 2022 M5.8 Cook Strait earthquake. *New Zealand Journal of Geology and Geophysics*, 68(1), 135–150. <https://doi.org/10.1080/00288306.2023.2284276>
- Chen, D.-Y., Wu, Y.-M., & Chin, T.-L. (2015). Incorporating low-cost seismometers into the central weather bureau seismic network for earthquake early warning in Taiwan. *Terrestrial, Atmospheric & Oceanic Sciences*, 26(5).
- Chung, A. I., Henson, I., & Allen, R. M. (2019). Optimizing earthquake early warning performance: Elarms-3. *Seismological Research Letters*, 90(2A), 727–743.
- Cochran, E. S., & Husker, A. L. (2019). How low should we go when warning for earthquakes?. *Science*, 366(6468), 957-958.
- Cochran, E. S., Bunn, J., Minson, S. E., Baltay, A. S., Kilb, D. L., Kodera, Y., & Hoshiba, M. (2019). Event detection performance of the plum earthquake early warning algorithm in southern California. *Bulletin of the Seismological Society of America*, 109(4), 1524–1541.
- Cochran, E. S., Saunders, J. K., Minson, S. E., Bunn, J., Baltay, A., Kilb, D., ... & Kodera, Y. (2022). Alert optimization of the PLUM earthquake early warning algorithm for the western United States. *Bulletin of the Seismological Society of America*, 112(2), 803-819.
- Cousins, J., & McVerry, G. H. (2010). Overview of strong-motion data from the Darfield earthquake. *Bulletin of the New Zealand Society for Earthquake Engineering*, 43(4), 222-227.

- Coutant, O. (1989). Program of numerical simulation AXITRA. *Res. Rep. LGIT (in French), Universite Joseph Fourier, Grenoble.*
- Cremen, G., & Galasso, C. (2020). Earthquake early warning: Recent advances and perspectives. *Earth-Science Reviews, 205*, 103184. <https://doi.org/10.1016/j.earscirev.2020.103184>
- Cremen, G., Galasso, C., & Zuccolo, E. (2022). Investigating the potential effectiveness of earthquake early warning across Europe. *Nature communications, 13*(1), 639.
- Crempien, J. G., & Archuleta, R. J. (2015). UCSB method for simulation of broadband ground motion from kinematic earthquake sources. *Seismological Research Letters, 86*(1), 61-67.
- Crowell, B. W. (2024). Earthquake and tsunami early warning with GNSS data. In *GNSS Monitoring of the Terrestrial Environment* (pp. 111-127). Elsevier.
- Crowell, B. W., Schmidt, D. A., Bodin, P., Vidale, J. E., Baker, B., Barrientos, S., & Geng, J. (2018). G-fast earthquake early warning potential for great earthquakes in Chile. *Seismological Research Letters, 89*(2A), 542–556.
- Cubrinovski, M., Bradley, B. A., Elwood, K. J., Johnston, D., Orchiston, C., Sullivan, T., & Wotherspoon, L. M. (2020). Wellington’s earthquake resilience: Lessons from the 2016 Kaikōura earthquake. *Earthquake Spectra, 36*(3), 1448–1484. <https://doi.org/10.1193/120317EQS349M>
- Cuellar, A., Espinosa-Aranda, J., Suarez, R., Ibarrola, G., Uribe, A., Rodríguez, F., Frontana, B. (2014). The Mexican seismic alert system (SASMEX): Its alert signals, broadcast results and performance during the m 7.4 Punta Maldonado earthquake of March 20th, 2012. In *Early warning for geological disasters* (pp. 71–87). Springer.
- Cultrera, G., Ameri, G., Saraò, A., Cirella, A., & Emolo, A. (2013). Ground-motion simulations within ShakeMap methodology: application to the 2008 Iwate-Miyagi Nairiku (Japan) and 1980 Irpinia (Italy) earthquakes. *Geophysical Journal International, 193*(1), 220-237.
- Dhakal, Y. P., & Kunugi, T. (2021). An evaluation of strong-motion parameters at the S-net ocean-bottom seismograph sites near the Kanto basin for earthquake early warning. *Frontiers in Earth Science, 9*, 699439.
- Dhakal, Y. P., Aoi, S., Kunugi, T., Suzuki, W., & Kimura, T. (2017). Assessment of nonlinear site response at ocean bottom seismograph sites based on S-wave horizontal-to-vertical spectral ratios: a study at the Sagami Bay area K-NET sites in Japan. *Earth, Planets and Space, 69*, 1-7.
- Donahue, J. L., & Abrahamson, N. A. (2014). Simulation-based hanging wall effects. *Earthquake Spectra, 30*(3), 1269-1284.
- Dowrick, D. (1996). The Modified Mercalli earthquake intensity scale. *Bulletin of the New Zealand Society for Earthquake Engineering, 29*(2), 92–106.
- Ely, G. P., Day, S. M., & Minster, J. B. (2009). A support-operator method for 3-D rupture dynamics. *Geophysical Journal International, 177*(3), 1140-1150.
- Farghal, N. S., Saunders, J. K., & Parker, G. A. (2022). The potential of using fiber optic distributed acoustic sensing (DAS) in earthquake early warning applications. *Bulletin of the Seismological Society of America, 112*(3), 1416-1435.

- Frankel, A., Wirth, E., Marafi, N., Vidale, J., & Stephenson, W. (2018). Broadband synthetic seismograms for magnitude 9 earthquakes on the Cascadia megathrust based on 3D simulations and stochastic synthetics, Part 1: Methodology and overall results. *Bulletin of the Seismological Society of America*, 108(5A), 2347-2369.
- Fujinawa, Y., & Noda, Y. (2013). Japan's earthquake early warning system on 11 March 2011: performance, shortcomings, and changes. *Earthquake Spectra*, 29(1\_suppl), 341-368.
- GeoNet. (2023, November 13). *GeoNet and Geohazard Monitoring in Aotearoa: What we do, how we do it, and a peek behind the scenes* [Workshop]. New Zealand.
- Gerstenberger, M. C., Bora, S., Bradley, B. A., DiCaprio, C., Kaiser, A., Manea, E. F., ... & Wotherspoon, L. M. (2024). The 2022 Aotearoa New Zealand national seismic hazard model: Process, overview, and results. *Bulletin of the Seismological Society of America*, 114(1), 7–36.
- Gerstenberger, M. C., Van Dissen, R., Rollins, C., DiCaprio, C., Thingbaijim, K. K., Bora, S., ... & Williams, C. (2024). The seismicity rate model for the 2022 Aotearoa New Zealand national seismic hazard model. *Bulletin of the Seismological Society of America*, 114(1), 182-216.
- GFZ. (2021). *Seedlink*. Retrieved from <https://www.seiscomp.de/doc/apps/seedlink.html>
- Gledhill, K., Ristau, J., Reyners, M., Fry, B., & Holden, C. (2010). The Darfield (Canterbury) earthquake of September 2010: preliminary seismological report. *Bulletin of the New Zealand Society for Earthquake Engineering*, 43(4), 215-221.
- GNS Science. (2017). *Rupture-models*. Retrieved from [https://github.com/GeoNet/data/blob/main/rupture-models/20100904-darfield/3366146\\_seismic\\_model.txt](https://github.com/GeoNet/data/blob/main/rupture-models/20100904-darfield/3366146_seismic_model.txt)
- GNS Science. (2022). *GeoNet Strong Motion Data Products [Dataset]*. GNS Science, GeoNet <https://doi.org/10.21420/X0MD-MV58>
- GNS Science. (2023a). *Earthquake intensity*. Retrieved from <https://www.geonet.org.nz/earthquake/intensity>
- GNS Science. (2023b). *GeoNet Aotearoa New Zealand Stations Metadata Repository [Dataset]*. GNS Science, GeoNet. <https://doi.org/10.21420/0VY2-C144>
- GNS Science. (2024a). *FDSN webservice*. Retrieved from <https://www.geonet.org.nz/data/tools/FDSN>
- GNS Science. (2024b). *GeoNet quakes - weak and above*. Retrieved from <https://www.geonet.org.nz/earthquake>
- GNS Science. (2024c). *Map and search of GeoNet sensor networks*. Retrieved from <https://www.geonet.org.nz/data/network/sensor/search>
- GNS Science. (n.d.). Darfield earthquake. <https://www.geonet.org.nz/earthquake/felt/3366146>
- Goltz, J. D., Wald, D. J., McBride, S. K., Reddy, E., Quitoriano, V., & Saunders, J. K. (2024). The Ojai California Earthquake of 20 August 2023: Earthquake Early Warning Performance and Alert Recipient Response in the M w 5.1 Event. *Seismological Research Letters*, 95(5), 2745-2760.
- Grapenthin, R., Johanson, I. A., & Allen, R. M. (2014b). Operational real-time GPS-enhanced earthquake early warning. *Journal of Geophysical Research: Solid Earth*, 119(10), 7944–7965.

- Grapenthin, R., Johanson, I., & Allen, R. M. (2014a). The 2014 Mw 6.0 Napa earthquake, California: Observations from real-time GPS-enhanced earthquake early warning. *Geophysical Research Letters*, *41*(23), 8269-8276.
- Greater Wellington Regional Council. (2024). *Greater Wellington Region Basemap*. ArcGIS Open Data. Retrieved from <https://data-gwrc.opendata.arcgis.com/maps/GWRC::greater-wellington-region-basemap/explore>
- Gusman, A. R., Roger, J., Power, W., Fry, B., & Kaneko, Y. (2022). The 2021 Loyalty Islands earthquake (Mw 7.7): Tsunami waveform inversion and implications for tsunami forecasting for New Zealand. *Earth and Space Science*, *9*(11), e2022EA002346.
- He, X., & Pan, H. (2024). A modified stochastic finite-fault method for estimating strong ground motion: Validation and application. *Earthquake Science*, *37*(1), 36-50.
- Hicks, S. P., & Rietbrock, A. (2015). Seismic slip on an upper-plate normal fault during a large subduction megathrust rupture. *Nature Geoscience*, *8*(12), 955-960.
- Hill, M. (2016). Anniversary of the 1931 Hawke's Bay earthquake. *Australian Journal of Emergency Management, The*, *31*(2), 58–59.
- Holden, C., & Kaiser, A. (2016). Stochastic ground motion modelling of the largest Mw 5.9+ aftershocks of the Canterbury 2010–2011 earthquake sequence. *New Zealand Journal of Geology and Geophysics*, *59*(1), 187-201.
- Holden, C., Beavan, J., Fry, B., Reyners, M., Ristau, J., Van Dissen, R., ... & Quigley, M. (2011, April). Preliminary source model of the Mw 7.1 Darfield earthquake from geological, geodetic and seismic data. In *9th Pacific conference on earthquake engineering* (pp. 1063-1092).
- Holden, C., Zhao, J., & Stirling, M. (2013). Ground motion modelling of a large subduction interface earthquake in Wellington, New Zealand. In *NZSEE Annual Conference*.
- Hoshiaba, M. (2013). Real-time prediction of ground motion by Kirchhoff Fresnel boundary integral equation method: Extended front detection method for earthquake early warning. *Journal of Geophysical Research: Solid Earth*, *118*(3), 1038–1050.
- Hoshiaba, M., & Aoki, S. (2015). Numerical shake prediction for earthquake early warning: Data assimilation, real-time shake mapping, and simulation of wave propagation. *Bulletin of the Seismological Society of America*, *105*(3), 1324–1338.
- Hoshiaba, M., Ohtake, K., Iwakiri, K., Aketagawa, T., Nakamura, H., & Yamamoto, S. (2010). How precisely can we anticipate seismic intensities? a study of uncertainty of anticipated seismic intensities for the earthquake early warning method in Japan. *Earth, planets and space*, *62*(8), 611–620.
- Hsiao, N.-C., Lin, T.-W., Hsu, S.-K., Kuo, K.-W., Shin, T.-C., & Leu, P.-L. (2014). Improvement of earthquake locations with the marine cable hosted observatory (macho) offshore ne Taiwan. *Marine Geophysical Research*, *35*(3), 327–336.
- Hsu, T.-Y., Wang, H.-H., Lin, P.-Y., Lin, C.-M., Kuo, C.-H., & Wen, K.-L. (2016). Performance of the NCREE's on-site warning system during the 5 February 2016 Mw 6.53 Meinong earthquake. *Geophysical Research Letters*, *43*(17), 8954–8959.

- Insurance Council of New Zealand. (n.d.). *Canterbury Earthquakes*. Retrieved, from <https://www.icnz.org.nz/industry/canterbury-earthquakes/>
- IRIS. (2021). *Miniseed*. Incorporated Research Institutions for Seismology. Retrieved from <http://ds.iris.edu/ds/nodes/dmc/data/formats/miniseed>
- Iwakiri, K., Hoshiba, M., Nakamura, K., & Morikawa, N. (2011). Improvement in the accuracy of expected seismic intensities for earthquake early warning in Japan using empirically estimated site amplification factors. *Earth, planets and space*, *63*, 57-69.
- JENLOGIX. (2021). *Palert station status*. Retrieved from <http://203.167.219.225/index>
- Kaiser, A. E., Hill, M. P., de la Torre, C., Bora, S., Manea, E., Wotherspoon, L., ... & Gerstenberger, M. (2024). Overview of site effects and the application of the 2022 New Zealand NSHM in the Wellington Basin, New Zealand. *Bulletin of the Seismological Society of America*, *114*(1), 399–421.
- Kaiser, A., Balfour, N., Fry, B., Holden, C., Litchfield, N., Gerstenberger, M., ... & Gledhill, K. (2017a). The 2016 Kaikōura, New Zealand, earthquake: Preliminary seismological report. *Seismological Research Letters*, *88*(3), 727–739. <https://doi.org/10.1785/0220170039>
- Kaiser, A., Van Houtte, C., Perrin, N., Wotherspoon, L., & McVerry, G. (2017b). Site characterisation of geonet stations for the New Zealand strong motion database. *Bulletin of the New Zealand Society for Earthquake Engineering*, *50*(1), 39–49.
- Kanamori, H. (2005). Real-time seismology and earthquake damage mitigation. *Annu. Rev. Earth Planet. Sci.*, *33*(1), 195-214.
- Kawamoto, S., Ohta, Y., Hiyama, Y., Todoriki, M., Nishimura, T., Furuya, T., ... Miyagawa, K. (2017). Regard: A new GNSS-based real-time finite fault modeling system for geonet. *Journal of Geophysical Research: Solid Earth*, *122*(2), 1324–1349.
- Kilb, D., Bunn, J., Saunders, J., Cochran, E., Minson, S., Baltay, A., ... Kodera, Y. (2021). The PLUM Earthquake early warning Algorithm: A retrospective case study of West Coast, USA, data. *Journal of Geophysical Research: Solid Earth*, *126*(7), e2020JB021053.
- Kodera, Y. (2019). An earthquake early warning method based on Huygens principle: Robust ground motion prediction using various localized distance-attenuation models. *Journal of Geophysical Research: Solid Earth*, *124*(12), 12981–12996.
- Kodera, Y., Hayashimoto, N., Moriwaki, K., Noguchi, K., Saito, J., Akutagawa, J., Adachi, S., Morimoto, M., Okamoto, K., Honda, S., & Hoshiba, M. (2020). First-Year Performance of a Nationwide Earthquake Early Warning System Using a Wavefield-Based Ground-Motion Prediction Algorithm in Japan. *Seismological Research Letters*, *91*(2A), 826–834.
- Kodera, Y., Hayashimoto, N., Tamaribuchi, K., Noguchi, K., Moriwaki, K., Takahashi, R., ... Hoshiba, M. (2021). Developments of the nationwide earthquake early warning system in Japan after the 2011 Mw 9.0 Tohoku-Oki earthquake. *Frontiers in Earth Science*, *9*, 726045. <https://doi.org/10.3389/feart.2021.726045>
- Kodera, Y., Yamada, Y., Hirano, K., Tamaribuchi, K., Adachi, S., Hayashimoto, N., ... Hoshiba, M. (2018). The propagation of local undamped motion (plum) method: A simple and robust seismic wavefield estimation approach for earthquake early warning. *Bulletin of the Seismological Society of America*, *108*(2), 983–1003.

- Kohler, M. D., Cochran, E. S., Given, D., Guiwits, S., Neuhauser, D., Henson, I., . . . Schwarz, S. (2017). Earthquake Early Warning ShakeAlert System: West Coast Wide production Prototype. *Seismological Research Letters*, *89*(1), 99–107. <https://doi.org/10.1785/0220170140>
- Kohler, M. D., Smith, D. E., Andrews, J., Chung, A. I., Hartog, R., Henson, I., ... & Guiwits, S. (2020). Earthquake early warning ShakeAlert 2.0: Public rollout. *Seismological Research Letters*, *91*(3), 1763-1775.
- Kohler, M. D., Smith, D. E., Andrews, J., Chung, A. I., Hartog, R., Henson, I., ... & Guiwits, S. (2020). Earthquake early warning ShakeAlert 2.0: Public rollout. *Seismological Research Letters*, *91*(3), 1763-1775.
- Kolb, A., & Conitzer, V. (2020). Crying about a strategic wolf: A theory of crime and warning. *Journal of Economic Theory*, *189*, 105094.
- Komatitsch, D., & Tromp, J. (1999). Introduction to the spectral element method for three-dimensional seismic wave propagation. *Geophysical journal international*, *139*(3), 806-822.
- Kuyuk, H. S., & Allen, R. M. (2013). Optimal seismic network density for earthquake early warning: A case study from California. *Seismological Research Letters*, *84*(6), 946-954.
- Lacoua, L., Fry, B., Gorman, A., Mika Liao, Y. W., Foundotos, L., Zweck, C., & Jamelot, A. (2024, April). Improving Rapid Earthquake Characterization of Tsunami Early Warning for Aotearoa New Zealand and the Southwest Pacific. In *EGU General Assembly Conference Abstracts* (p. 479).
- Lancieri, M., & Zollo, A. (2008). A Bayesian approach to the real-time estimation of magnitude from the early p and s wave displacement peaks. *Journal of Geophysical Research: Solid Earth*, *113*(B12).
- Lux, A. I., Smith, D., Böse, M., McGuire, J. J., Saunders, J. K., Huynh, M., ... Toomey, D. (2024). Status and performance of the ShakeAlert earthquake early warning system: 2019–2023. *Bulletin of the Seismological Society of America*, *114*(6), 3041-3062. <https://doi.org/10.1785/0120230453>
- Mackie, B. (2014). *Warning fatigue: Insights from the Australian bushfire context* (Doctoral dissertation). University of Canterbury, New Zealand.
- Mai, P. M., Imperatori, W., & Olsen, K. B. (2010). Hybrid broadband ground-motion simulations: Combining long-period deterministic synthetics with high-frequency multiple s-to-s backscattering. *Bulletin of the Seismological Society of America*, *100*(5A), 2124–2142.
- Massin, F., Clinton, J., & Böse, M. (2021). Status of earthquake early warning in Switzerland. *Frontiers in Earth Science*, *9*, 707654.
- Mateparae, J. (2015). National civil defence emergency management plan order 2015. *Published under the authority of the New Zealand Government, Wellington, New Zealand*.
- McBride, S. K., Sumy, D. F., Llenos, A. L., Parker, G. A., McGuire, J., Saunders, J. K., Meier, M-A., Shuback, P., Given, D. & de Groot, R. (2023). Latency and geofence testing of wireless emergency alerts intended for the ShakeAlert® earthquake early warning system for the West Coast of the United States of America. *Safety science*, *157*, 105898.
- Meier, M. A. (2017). How “good” are real-time ground motion predictions from earthquake early warning systems?. *Journal of Geophysical Research: Solid Earth*, *122*(7), 5561-5577.

- Meier, M. A., Kodera, Y., Böse, M., Chung, A., Hoshiya, M., Cochran, E., ... & Heaton, T. (2020). How often can earthquake early warning systems alert sites with high-intensity ground motion?. *Journal of Geophysical Research: Solid Earth*, *125*(2), e2019JB017718.
- Minson, S. E., Baltay, A. S., Cochran, E. S., Hanks, T. C., Page, M. T., McBride, S. K., ... & Meier, M. A. (2019). The limits of earthquake early warning accuracy and best alerting strategy. *Scientific reports*, *9*(1), 1-13.
- Minson, S. E., Cochran, E. S., Saunders, J. K., McBride, S. K., Wu, S., Baltay, A. S., & Milner, K. R. (2022). What to expect when you are expecting earthquake early warning. *Geophysical Journal International*, *231*(2), 1386-1403.
- Minson, S. E., Murray, J. R., Langbein, J. O., & Gomberg, J. S. (2014). Real-time inversions for finite fault slip models and rupture geometry based on high-rate GPS data. *Journal of Geophysical Research: Solid Earth*, *119*(4), 3201–3231.
- Minson, S. E., Saunders, J. K., Bunn, J. J., Cochran, E. S., Baltay, A. S., Kilb, D. L., ... Kodera, Y. (2020). Real-time performance of the plum earthquake early warning method during the 2019 m 6.4 and 7.1 Ridgecrest, California, earthquakes. *Bulletin of the Seismological Society of America*, *110*(4), 1887–1903. <https://doi.org/10.1785/0120200038>
- Moratalla, J. M., Goded, T., Rhoades, D. A., Canessa, S., & Gerstenberger, M. C. (2021). New ground motion to intensity conversion equations (GMICEs) for New Zealand. *Seismological Society of America*, *92*(1), 448-459.
- Motazedian, D., & Atkinson, G. M. (2005). Stochastic finite-fault modeling based on a dynamic corner frequency. *Bulletin of the Seismological Society of America*, *95*(3), 995-1010.
- Nakayachi, K., Becker, J. S., Potter, S. H., & Dixon, M. (2019). Residents' reactions to earthquake early warnings in Japan. *Risk Analysis*, *39*(8), 1723-1740.
- National Earthquake Hazards Reduction Program (NEHRP). (2023). *Seismic site classification*. Retrieved from <https://www.masw.com/files/NEHRP.pdf>
- NEC. (2021). *NEC delivers submarine cable seismic and tsunami observation system to Taiwan's central weather bureau*. Retrieved from <https://www.nec.com/en/press/202101/global2021012601.html>
- NEMA. (2020a). *Tsunami warnings-media-guide*. Retrieved from <https://www.civildefence.govt.nz/assets/Uploads/media/Guides-for-media/media-guide-tsunami-warnings-4-sep-2020.pdf>
- NEMA. (2020b). *National tsunami advisory and warning plan supporting plan [sp 01/20] revised June 2020*. Retrieved from <https://www.civildefence.govt.nz/cdem/sector/guidelines/national-tsunami-advisory-and-warning-plan>
- NEMA. (2021). *Emergency mobile alert*. Retrieved from <https://www.civildefence.govt.nz/get-ready/civil-defence-emergency>
- Nishimae, Y., Kodera, Y., Yamada, Y., Adachi, S., Morimoto, M., & Hoshiya, M. (2016). Review of Earthquake Early Warning Operation in Japan for eight years from 2007. *Japan Geoscience Union*.

- Orense, R. P., Mirjafari, Y., Asadi, S., Naghibi, M., Chen, X., Altaf, O., & Asadi, B. (2017). Ground performance in Wellington waterfront area following the 2016 Kaikōura earthquake. *Bulletin of the New Zealand Society for Earthquake Engineering*, 50(2), 142-151.
- Oth, A., Böse, M., Wenzel, F., Köhler, N., & Erdik, M. (2010). Evaluation and optimization of seismic networks and algorithms for earthquake early warning—the case of Istanbul (Turkey). *Journal of Geophysical Research: Solid Earth*, 115(B10).
- Pacific Earthquake Engineering Research Center. (2023). *PEER Ground Motion Database*. University of California, Berkeley. Retrieved from <https://ngawest2.berkeley.edu>
- Pagani, M., Monelli, D., Weatherill, G., Danciu, L., Crowley, H., Silva, V., ... & Vigano, D. (2014). OpenQuake engine: An open hazard (and risk) software for the global earthquake model. *Seismological Research Letters*, 85(3), 692-702.
- Papadopoulos, A. N., Böse, M., Danciu, L., Clinton, J., & Wiemer, S. (2023). A framework to quantify the effectiveness of earthquake early warning in mitigating seismic risk. *Earthquake Spectra*, 39(2), 938–961. <https://doi.org/10.1193/220216EQS302M>
- Peng, H., Wu, S., & Yamada, M. (2023, December). Constructing an empirical envelope function of seismic waveforms for the evaluation of EEW in Japan. In *AGU Fall Meeting Abstracts* (Vol. 2023, No. 7, pp. NH13C-07).
- Perrin, N. D., Heron, D., Kaiser, A., & Van Houtte, C. (2015, April). VS30 and NZS 1170.5 site class maps of New Zealand. In *New Zealand Society for Earthquake Engineering Annual Technical Conference* (pp. 10-12).
- Peter, D., Komatitsch, D., Luo, Y., Martin, R., Le Goff, N., Casarotti, E., ... & Tromp, J. (2011). Forward and adjoint simulations of seismic wave propagation on fully unstructured hexahedral meshes. *Geophysical Journal International*, 186(2), 721-739.
- Petersen, M. D., Moschetti, M. P., Powers, P. M., Mueller, C. S., Haller, K. M., Frankel, A. D., ... & Olsen, A. H. (2015). The 2014 United States national seismic hazard model. *Earthquake Spectra*, 31(1\_suppl), S1-S30.
- Petersen, T., Gledhill, K., Chadwick, M., Gale, N. H., & Ristau, J. (2011). The New Zealand national seismograph network. *Seismological Research Letters*, 82(1), 9–20.
- Porras, J., Massin, F., Arroyo-Solórzano, M., Arroyo, I., Linkimer, L., Böse, M., & Clinton, J. (2021). Preliminary results of an earthquake early warning system in Costa Rica. *Frontiers in Earth Science*, 9, 700843.
- Potter, S. H., Becker, J. S., Johnston, D. M., & Rossiter, K. P. (2015). An overview of the impacts of the 2010-2011 canterbury earthquakes. *International Journal of Disaster Risk Reduction*, 14, 6–14.
- Prasanna, R., Chandrakumar, C., Nandana, R., Holden, C., Punchihewa, A., Becker, J. S., ... & Tan, M. L. (2022). “Saving Precious Seconds”—A novel approach to implementing a low-cost earthquake early warning system with node-level detection and alert generation. In *Informatics* (Vol. 9, No. 1, p. 25). MDPI.
- Quigley, M. C., Hughes, M. W., Bradley, B. A., van Ballegooy, S., Reid, C., Morgenroth, J., ... Pettinga, J. R. (2016). The 2010–2011 Canterbury earthquake sequence: Environmental effects, seismic triggering thresholds and geologic legacy. *Tectonophysics*, 672, 228-274. doi:10.1016/j.tecto.2016.01.044

- Rawlinson, Z. J., Townend, J., Arnold, R., & Bannister, S. (2012). Derivation and implementation of a nonlinear experimental design criterion and its application to seismic network expansion at Kawerau geothermal field, New Zealand. *Geophysical Journal International*, 191(2), 686-694.
- REANNZ. (2021). *Tracking New Zealand's movements*. Research and Education Advanced Network New Zealand. Retrieved from <https://www.reannz.co.nz/case-studies/tracking-new-zealands-movements>
- Richards-Dinger, K., & Dieterich, J. H. (2012). RSQSIM earthquake simulator. *Seismological Research Letters*, 83(6), 983-990.
- Ristau, J. (2008). Implementation of routine regional moment tensor analysis in New Zealand. *Seismological Research Letters*, 79(3), 400-415.
- Ruhl, C. J., Melgar, D., Chung, A. I., Grapenthin, R., & Allen, R. M. (2019). Quantifying the value of real-time geodetic constraints for earthquake early warning using a global seismic and geodetic data set. *Journal of Geophysical Research: Solid Earth*, 124(4), 3819-3837.
- Ruhl, C. J., Melgar, D., Chung, A. I., Grapenthin, R., & Allen, R. M. (2019). Quantifying the value of real-time geodetic constraints for earthquake early warning using a global seismic and geodetic data set. *Journal of Geophysical Research: Solid Earth*, 124(4), 3819-3837.
- Ruhl, C. J., Melgar, D., Grapenthin, R., & Allen, R. M. (2017). The value of real-time GNSS to earthquake early warning. *Geophysical Research Letters*, 44(16), 8311-8319.
- Ruhl, C. J., Melgar, D., Grapenthin, R., & Allen, R. M. (2017). The value of real-time GNSS to earthquake early warning. *Geophysical Research Letters*, 44(16), 8311-8319.
- Rydelek, P., Wu, C., & Horiuchi, S. (2007). Comment on "Earthquake magnitude estimation from peak amplitudes of very early seismic signals on strong motion records" by Aldo Zollo, Maria Lancieri, and Stefan Nielsen. *Geophysical Research Letters*, 34(20).
- Satriano, C., Elia, L., Martino, C., Lancieri, M., Zollo, A., & Lannaccone, G. (2011). Presto, the earthquake early warning system for southern Italy: Concepts, capabilities, and future perspectives. *Soil Dynamics and Earthquake Engineering*, 31(2), 137-153.
- Satriano, C., Lomax, A., & Zollo, A. (2008). Real-time evolutionary earthquake location for seismic early warning. *Bulletin of the Seismological Society of America*, 98(3), 1482-1494.
- Satriano, C., Wu, Y.-M., Zollo, A., & Kanamori, H. (2011). Earthquake early warning: Concepts, methods, and physical grounds. *Soil Dynamics and Earthquake Engineering*, 31(2), 106-118.
- Saunders, J. K., Cochran, E. S., Bunn, J. J., Baltay, A. S., Minson, S. E., & O'Rourke, C. T. (2024). Incorporating intensity distance attenuation into PLUM ground-motion-based earthquake early warning in the United States: The APPLES configuration. *Earth's Future*, 12(2), e2023EF004126.
- Saunders, J. K., Minson, S. E., & Baltay, A. S. (2022a). How low should we alert? Quantifying intensity threshold alerting strategies for earthquake early warning in the United States. *Earth's Future*, 10(3).
- Saunders, J. K., Minson, S. E., Baltay, A. S., Bunn, J. J., Cochran, E. S., Kilb, D. L., ... & Kodera, Y. (2022b). Real-time earthquake detection and alerting behavior of PLUM ground-motion-based early warning in the United States. *Bulletin of the Seismological Society of America*, 112(5), 2668-2688.

- Serdar Kuyuk, H., Allen, R. M., Brown, H., Hellweg, M., Henson, I., & Neuhauser, D. (2014). Designing a network-based earthquake early warning algorithm for California: Elarms-2. *Bulletin of the Seismological Society of America*, 104(1), 162–173.
- ShakeAlert. (2024, December). *Alert delivery thresholds*. Retrieved December 2024, from <https://www.shakealert.org/system-information/alert-delivery-thresholds/>
- Shaw, B. E. (2023). Magnitude and slip scaling relations for fault-based seismic hazard. *Bulletin of the Seismological Society of America*, 113(3), 924-947.
- Shaw, B. E., Fry, B., Nicol, A., Howell, A., & Gerstenberger, M. (2022). An earthquake simulator for New Zealand. *Bulletin of the Seismological Society of America*, 112(2), 763-778.
- Sheen, D. H., Park, J. H., Chi, H. C., Hwang, E. H., Lim, I. S., Seong, Y. J., & Pak, J. (2017). The first stage of an earthquake early warning system in South Korea. *Seismological Research Letters*, 88(6), 1491-1498. <https://doi.org/10.1785/0220170133>
- Spooner, B. (2021). *Introducing android earthquake alerts outside the U.S.* Google. Retrieved from <https://blog.google/products/android/introducing-android-earthquake-alerts-outside-us>
- Spudich, P., & Xu, L. (2002). Documentation of software package COMPSYN svx3. 11: Programs for earthquake ground motion calculation using complete 1-d green's functions. *International Handbook of Earthquake and Engineering Seismology*, 26, 64.
- Stevens, C., & Smith, M. (2009). Cook Strait: the nexus in New Zealand's coastal oceanography. In *coasts and ports 2009: In a dynamic environment* (pp. 80-86). [Wellington, NZ]: Engineers Australia.
- Stirling, M., McVerry, G., Gerstenberger, M., Litchfield, N., Van Dissen, R., Berryman, K., ... & Jacobs, K. (2012). National seismic hazard model for New Zealand: 2010 update. *Bulletin of the Seismological Society of America*, 102(4), 1514-1542.
- Strauss, J. A., & Allen, R. M. (2016). Benefits and costs of earthquake early warning. *Seismological Research Letters*, 87(3), 765–772. <https://doi.org/10.1785/0220150215>
- Stuart-Black, S. (2021). New technology boosting early warning capability in New Zealand. *Australian Journal of Emergency Management*, 36(1), 26–27.
- Stubailo, I., Alvarez, M., Biasi, G., Bhadha, R., & Hauksson, E. (2020). Latency of waveform data delivery from the Southern California seismic network during the 2019 Ridgecrest earthquake sequence and its effect on ShakeAlert. *Seismological Society of America*, 92(1), 170-186. <https://doi.org/10.1785/0220200211>
- Suárez, G. (2022). The seismic early warning system of Mexico (SASMEX): A retrospective view and future challenges. *Frontiers in Earth Science*, 10, 827236. <https://doi.org/10.3389/feart.2022.827236>
- Suárez, G., van Eck, T., Giardini, D., Ahern, T., Butler, R., & Tsuboi, S. (2008). The international federation of digital seismograph networks (FDSN): An integrated system of seismological observatories. *IEEE Systems Journal*, 2(3), 431-438.
- Sukhovich, A., Bonnieux, S., Hello, Y., Irissou, J. O., Simons, F. J., & Nolet, G. (2015). Seismic monitoring in the oceans by autonomous floats. *Nature communications*, 6(1), 8027.

- Taborda, R., López, J., Karaoglu, H., Urbanic, J., & Bielak, J. (2010). Speeding up finite element wave propagation for large-scale earthquake simulations. *Parallel Data Laboratory Tech. Rept. CMUPDL, 10*, 109.
- Tajima, F., & Hayashida, T. (2018). Earthquake early warning: what does “seconds before a strong hit” mean? *Progress in Earth and Planetary Science, 5*(1), 1–25.
- Tan, M. L., Leonard, G. S., & Johnston, D. M. (2021). *Hawke’s Bay Regional Alerting Systems Review* (Disaster Research Science Report 2021:4, pp. 1–50). Joint Centre for Disaster Research – Massey University.
- Tan, M. L., Vinnell, L. J., Valentin, A. P. M., Prasanna, R., & Becker, J. S. (2023). The public's perception of an earthquake early warning system: A study on factors influencing continuance intention. *International journal of disaster risk reduction, 97*, 104032.
- Townend, J., Behr, Y., Buckley, K., Savage, M., & Hine, J. (2009). A grid-based facility for large-scale cross-correlation of continuous seismic data. *EQC*.
- Townend, J., Villamor, P., & Quigley, M. (2012). Introduction to the Canterbury earthquake sequence special issue. *New Zealand Journal of Geology and Geophysics, 55*(3), 151-154.  
doi:10.1080/00288306.2012.712539
- Trugman, D. T., Page, M. T., Minson, S. E., & Cochran, E. S. (2019). Peak ground displacement saturates exactly when expected: Implications for earthquake early warning. *Journal of Geophysical Research: Solid Earth, 124*(5), 4642-4653.
- Tsuno, S., Niwa, K., Korenaga, M., Yamanaka, H., Chimoto, K., Miyake, H., & Yamada, N. (2024). Application of the on-site P-wave earthquake early warning method based on site-specific ratios of S-waves to P-waves to the 2016 Kumamoto earthquake sequence, Japan. *Earth, Planets and Space, 76*(1), 32.
- Tu, T., Yu, H., Ramirez-Guzman, L., Bielak, J., Ghattas, O., Ma, K. L., & O'hallaron, D. R. (2006, November). From mesh generation to scientific visualization: An end-to-end approach to parallel supercomputing. In *Proceedings of the 2006 ACM/IEEE conference on Supercomputing* (pp. 91-es).
- USGS. (2021a). *Entire U.S. West Coast Now Has Access to Shakealert®*. Retrieved from <https://www.usgs.gov/news/entire-us-west-coast-now-has-access-shakealert-earthquake-early-warning?qt-newsscienceproducts=4#qt-newsscienceproducts>
- USGS. (2021b). *Shakealert®*. Retrieved from <https://www.shakealert.org>
- Van Houtte, C., Bannister, S., Holden, C., Bourguignon, S., & McVerry, G. (2017). The New Zealand strong motion database. *Bulletin of the New Zealand Society for Earthquake Engineering, 50*(1), 1–20
- Vinnell, L. J., Hudson-Doyle, E. E., Inch, P., Tan, M. L., Becker, J. S., & Johnston, D. M. (2022). Evacuation behavior and information needs of Wellington, Aotearoa New Zealand residents following the 5 March 2021 Mw 7.3 East Cape earthquake. *Seismological Research Letters, 93*(3), 1452-1463.
- Vinnell, L. J., Tan, M. L., Prasanna, R., & Becker, J. S. (2023). Knowledge, perceptions, and behavioral responses to earthquake early warning in Aotearoa New Zealand. *Frontiers in Communication, 8*, 1229247.

- Wald, D. J., Quitoriano, V., Heaton, T. H., & Kanamori, H. (1999). Relationships between peak ground acceleration, peak ground velocity, and modified Mercalli intensity in California. *Earthquake spectra*, 15(3), 557–564.
- Wald, D. J., Worden, B. C., Quitoriano, V., & Pankow, K. L. (2005). *Shakemap manual: technical manual, user's guide, and software guide* (Tech. Rep.).
- Wallace, L. M., Townend, J., Stevens, C., Kellett, R. L., De Spuza, J., & Giorli, G. (2021). Chatham Islands Cabled Observatory Science Opportunities: Workshop 23-24 February 2021, Summary Report. Lower Hutt (NZ). *GNS Sci*, 70.
- Wells, D. L., & Coppersmith, K. J. (1994). New empirical relationships among magnitude, rupture length, rupture width, rupture area, and surface displacement. *Bulletin of the seismological Society of America*, 84(4), 974-1002.
- Worden, C. B., & Wald, D. J. (2016). ShakeMap manual online: Technical manual, user's guide, and software guide. *US Geological Survey*, 1-156.
- Worden, C. B., Gerstenberger, M. C., Rhoades, D. A., & Wald, D. J. (2012). Probabilistic relationships between ground-motion parameters and modified Mercalli intensity in California. *Bulletin of the Seismological Society of America*, 102(1), 204-221.
- Worden, C. B., Hearne, M., Wald, D. J., Pagani, M., Thompson, E. M., Verros, S. A., & Weatherill, G. (2017). Complimentary components of OpenQuake and ShakeMap. In *Proc. of the 16th World Conf. on Earthquake Engineering*.
- Worldatlas. (2021). *New Zealand's deadliest disasters ever*. Retrieved from <https://www.worldatlas.com/articles/new-zealand-s-deadliest-disasters-ever.html>
- WorldPop. (2020). New Zealand 2020 population dataset. WorldPop, University of Southampton. Retrieved from [https://data.worldpop.org/GIS/Population/Global\\_2000\\_2020/2020/NZL/nzl\\_ppp\\_2020.tif](https://data.worldpop.org/GIS/Population/Global_2000_2020/2020/NZL/nzl_ppp_2020.tif)
- Wotherspoon, L. M., Kaiser, A. E., Stolte, A. C., & Manea, E. F. (2024). Development of the site characterization database for the 2022 New Zealand National Seismic Hazard Model. *Seismological Research Letters*, 95(1), 214-225.
- Wu, Y. M., & Kanamori, H. (2005). Experiment on an onsite early warning method for the Taiwan early warning system. *Bulletin of the Seismological Society of America*, 95(1), 347-353.
- Wu, Y., Hsiao, N., Chin, T., Chen, D., Chan, Y., & Wang, K. (2013). Earthquake early warning system in Taiwan. *Encyclopedia of earthquake engineering*, 12.
- Wu, Y.-M., & Kanamori, H. (2005). Experiment on an onsite early warning method for the Taiwan early warning system. *Bulletin of the Seismological Society of America*, 95(1), 347– 353.
- Wu, Y.-M., & Kanamori, H. (2008). Development of an earthquake early warning system using real-time strong motion signals. *Sensors*, 8(1), 1–9.
- Wu, Y.-M., Chen, D.-Y., Lin, T.-L., Hsieh, C.-Y., Chin, T.-L., Chang, W.-Y., ... Ker, S.-H. (2013). A high-density seismic network for earthquake early warning in Taiwan based on low cost sensors. *Seismological Research Letters*, 84(6), 1048–1054.

- Wu, Y.-M., Mittal, H., Chen, D. Y., Hsu, T. Y., & Lin, P. Y. (2021). Earthquake early warning systems in Taiwan: Current status. *Journal of the Geological Society of India*, 97, 1525-1532.  
<https://doi.org/10.1007/s12594-021-1909-6>
- Wurman, G., Allen, R. M., & Lombard, P. (2007). Toward earthquake early warning in northern California. *Journal of Geophysical Research: Solid Earth*, 112(B8).
- Xu, Y., Wang, J., Wu, Y.-M., & Kuo-Chen, H. (2017). Reliability assessment on earthquake early warning: A case study from Taiwan. *Soil dynamics and earthquake engineering*, 92, 397-407.
- Yamada, M., & Heaton, T. (2008). Real-time estimation of fault rupture extent using envelopes of acceleration. *Bulletin of the Seismological Society of America*, 98(2), 607-619.
- Yamamoto, S., Rydelek, P., Horiuchi, S., Wu, C., & Nakamura, H. (2008). On the estimation of seismic intensity in earthquake early warning systems. *Geophysical research letters*, 35(7).
- Yamasaki, E. (2012). What we can learn from Japan's earthquake early warning system. *Momentum*, 1-1.
- Zheng, X., Zhang, Y., Wang, R., Zhao, L., Li, W., & Huang, Q. (2020). Automatic inversions of strong-motion records for finite-fault models of significant earthquakes in and around Japan. *Journal of Geophysical Research: Solid Earth*, 125(9), e2020JB019992.
- Zollo, A., Colombelli, S., Elia, L., Emolo, A., Festa, G., Iannaccone, G., ... & Gasparini, P. (2014). An integrated regional and on-site Earthquake Early Warning System for Southern Italy: Concepts, methodologies and performances. *Early Warning for Geological Disasters: Scientific Methods and Current Practice*, 117-137.
- Zollo, A., Lancieri, M., & Nielsen, S. (2006). Earthquake magnitude estimation from peak amplitudes of very early seismic signals on strong motion records. *Geophysical Research Letters*, 33(23).